

Take One Step at a Time to Know Incremental Utility of Demonstration: An Analysis on Reranking for Few-Shot In-Context Learning

Kazuma Hashimoto Karthik Raman Michael Bendersky

Google Research, Mountain View

{kazumah, karthikraman, bemike}@google.com

Abstract

In-Context Learning (ICL) is an emergent capability of Large Language Models (LLMs). Only a few demonstrations enable LLMs to be used as blackbox for new tasks. Previous studies have shown that using LLMs’ outputs as labels is effective in training models to select demonstrations. Such a label is expected to estimate utility of a demonstration in ICL; however, it has not been well understood how different labeling strategies affect results on target tasks. This paper presents an analysis on different utility functions by focusing on LLMs’ output probability given ground-truth output, and task-specific reward given LLMs’ prediction. Unlike the previous work, we introduce a novel labeling method, incremental utility, which estimates how much incremental knowledge is brought into the LLMs by a demonstration. We conduct experiments with instruction-tuned LLMs on binary/multi-class classification, segmentation, and translation across Arabic, English, Finnish, Japanese, and Spanish. Our results show that (1) the probability is effective when the probability values are distributed across the whole value range (on the classification tasks), and (2) the downstream metric is more robust when nuanced reward values are provided with long outputs (on the segmentation and translation tasks). We then show that the proposed incremental utility further helps ICL by contrasting how the LLMs perform with and without the demonstrations.

1 Introduction

Recent advances of Large Language Models (LLMs) (Google et al., 2023; OpenAI, 2023; Touvron et al., 2023) have been pushing the field of Natural Language Processing (NLP) to the next level in many different aspects. A notable capability of LLMs is few-shot *In-Context Learning* (ICL), which uses only a few *demonstrations* (i.e.,

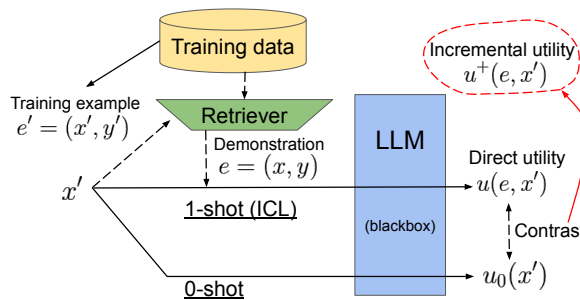


Figure 1: Incremental utility of a demonstration.

input-output pairs) to perform new tasks without finetuning (Brown et al., 2020; Zhao et al., 2021).

A crucial research topic is demonstration selection for ICL. Liu et al. (2022) proposed retrieval-based ICL, and Rubin et al. (2022) have shown that further finetuning retrieval models is effective (Li et al., 2023; Luo et al., 2023; Wang et al., 2023b). However, it is not conclusive how to train the selection models in several aspects: training labels, objectives, and overall pipelines.

This paper focuses on the training labels of the selection models. First, we use two types of values to estimate utility of a demonstration for an input: (1) an LLM’s output probability of generating the ground-truth output (as in the previous work), and (2) a task-specific reward function given the LLM’s prediction (as in reward optimization (Ranzato et al., 2016)). Next, we introduce a novel method to estimate *incremental utility* by contrasting the 1-shot ICL performance and the 0-shot performance (Figure 1). We expect the incremental utility to estimate how much incremental knowledge is brought by the demonstration, assuming that the LLMs have a capability of 0-shot prediction (Brown et al., 2020).

We conduct experiments in a standard retrieval-reranking framework with instruction-tuned LLMs. The evaluation tasks are binary/multi-class classification, segmentation, and translation across five

languages: Arabic, English, Finnish, Japanese, and Spanish. Our results and analysis show the following insights:

- the output probability of generating the ground-truth output is indeed good proxy for the utility of a demonstration, especially when the probability values are distributed across the whole value range with short outputs (e.g., classification),
- the downstream metric reward with nuanced values, on the other hand, is more robust for longer outputs (e.g., text generation tasks),
- our proposed incremental utility further improves ICL given contrastive training examples to effectively train the reranking model, and
- constrained retrieval is helpful when the retrieved candidates are imbalanced for training.

2 Inference and Training Pipelines

This section describes our experimental design for the comparative study. For a target task, we have a training set $\mathcal{T} = [e_1, \dots, e_i, \dots]$, where $e_i = (x'_i, y'_i)$ is an input-output pair; for example, y'_i is a class label on a text classification task.

2.1 Inference

Demonstration retrieval and reranking Given a new input x , we use a demonstration-selection model $M(e_i, x)$ to predict utility of a training example e_i as a demonstration. M is a cross-attention encoder that takes a concatenation of e_i and x as an input and then outputs a utility score. We select demonstrations with the k highest scores. To reduce the search space for the cross-attention encoder, we use an off-the-shelf text retriever to retrieve top- n training examples, according to similarity between x and x'_i . This is motivated by the success of the retrieval-based demonstration (Liu et al., 2022) and the standard retrieve-then-rerank framework (Zhuang et al., 2023). Hence we call M a *reranking* model.

ICL: prompting LLMs We create a prompt: $\text{prompt}(x, \mathcal{E}_x)$, where \mathcal{E}_x is the set of the top- k demonstrations selected from the n retrieved examples. The prompt is then fed into an LLM to generate a prediction as a string:

$$y = \text{LLM}(\text{prompt}(x, \mathcal{E}_x)), \quad (1)$$

where the prediction y is used for task-specific evaluation.

2.2 Training

We use a utility function $u(e_i, x)$ to train the reranking model M ; the larger the value is, the more e_i is expected to contribute to ICL for x . To create the training examples for M , we simulate the inference process within the training set; for each $x'_j \in \mathcal{T}$, we retrieve the top- n set and compute $u(e, x'_j)$ with $e \in \mathcal{E}_{x'_j}$. We ensure that e_j is not in $\mathcal{E}_{x'_j}$, and consequently we have $n|\mathcal{T}|$ values of $u(e, x)$ in total. Assuming that the value range of the utility and the reranking score is $[0.0, 1.0]$, we use logistic regression for pointwise training (Nie et al., 2019; Nogueira and Cho, 2019):

$$\begin{aligned} & -u(e_i, x) \log M(e_i, x) \\ & - (1 - u(e_i, x)) \log(1 - M(e_i, x)). \end{aligned} \quad (2)$$

3 Utility functions

We define multiple utility functions to investigate the effects on target tasks. Note that the utility functions are used only in the training phase.

3.1 Direct Utility

To compute $u(e, x'_j)$, it is considered to be the most straightforward to inspect how the LLM behaves given $\text{prompt}(x'_j, \{e\})$. We describe two utility functions that directly use the *1-shot* result.

Output Probability (OP) We follow a widely-used approach of using the likelihood of predicting the ground-truth output (Li et al., 2023):

$$u_{\text{OP}}(e, x'_j) = p(y'_j | \text{prompt}(x'_j, \{e\})), \quad (3)$$

assuming that we have access to the probability by feeding the prompt and y'_j together into the LLM.¹

Downstream Metric (DM) The primary goal is to improve downstream metrics of the target tasks. As in reward optimization (Ranzato et al., 2016), we define the following function:

$$u_{\text{DM}}(e, x'_j) = R(y^*, y'_j), \quad (4)$$

where $y^* = \text{LLM}(\text{prompt}(x'_j, \{e\}))$ is the LLM’s prediction, and R is a pre-defined reward function that is correlated with a task-specific metric;

¹Li et al. (2023) normalized the output probability values across possible output candidates for classification and multi-choice tasks. We do not apply the normalization, because the computational cost is not negligible when handling a large number of classes.

	$u_0(x'_j)$	$u(e, x'_j)$	$u(e, x'_j) - u_0(x'_j)$	$\max(u(e, x'_j), u_0(x'_j))$	$\ell = 0.0$	$\ell = 0.5$	$\ell = 0.8$	$\ell = 1.0$
(a)	0.0	0.1	0.1	0.1	0.1	0.316	0.631	1.0
(b)	0.9	1.0	0.1	1.0	0.1	0.1	0.1	0.1
(c)	0.0	0.0001	0.0001	0.0001	0.0001	0.01	0.158	1.0
(d)	0.3	0.5	0.2	0.5	0.2	0.283	0.348	0.4
(e)	0.5	0.3	-0.2	0.5	-0.2	-0.283	-0.348	-0.4

Table 1: Synthetic examples of $r(e, x'_j)$ in Equation (6) to illustrate how the incremental utility works.

	Train	Validation	Test
ISD (en)	3,122	346	1,400
ISD (ar)	2,792	310	1,400
EDOS-A (en)	14,000	2,000	4,000
EDOS-B (en)	3,398	486	970
CLINC (en)	15,100	3,100	5,500
SSENT (en)	1,744	249	499
SSENT (es)	1,438	206	410
XML-MT (ja)	100,033	500	1,500
XML-MT (fi)	97,893	500	1,500

Table 2: Dataset statistics.

Section 4.2 describes a reward function for each dataset.

3.2 Incremental Utility

Another way of interpreting the value of $u(e, x'_j)$ is how much *incremental* knowledge e brings into the LLM. For this, we would like to inspect the LLM’s capability of handling x'_j in the 0 -shot inference:

$$\text{LLM}(\text{prompt}(x'_j, \{\})), \quad (5)$$

where $\{\}$ represents an empty demonstration set. We can use the above-defined utility functions to evaluate the inference result; we use $u_0(x'_j)$ to denote a utility function without e .

For example, let us think about two cases where the use of e improves the value by 0.1: $(u_0(x'_j), u(e, x'_j)) = (0.9, 1.0)$ and $(0.0, 0.1)$. If we use $u(e, x'_j)$ alone, the former case has much larger value; however, this could underestimate the latter case’s value. We interpret those cases from different angles:

- (A) both have the same x'_j value in the absolute difference ($u(e, x'_j) - u_0(x'_j) = 0.1$), or
- (B) the absolute difference is only 10% of $u(e, x'_j)$ in the former case and 100% in the latter case.

The idea A has been explored in previous work (Li and Qiu, 2023), while the idea B is expected to be more reasonable with the example cases. One potential caveat is that the ratio-based approach overestimates the value, for example, with $(u_0(x'_j), u(e, x'_j)) = (0.0, 0.0001)$.

We propose incremental utility r to take into account all the aspects discussed above:

$$r(e, x'_j) = \frac{u(e, x'_j) - u_0(x'_j)}{\max(u(e, x'_j), u_0(x'_j))^\ell}, \quad (6)$$

where $\ell \in [0.0, 1.0]$ is a hyper-parameter. Here are the following key features we can read:

- $\ell = 0.0$ corresponds to the idea A, $\ell = 1.0$ to B, and the others take a balance between A and B; examples (a), (b), and (c) in Table 1 describe the above-mentioned cases.
- Examples (d) and (e) in Table 1 show the *symmetry* of the function.
- The value range of r is $[-1.0, 1.0]$, and $r < 0$ means that e *negatively* affects ICL.

Finally, we linearly transform r into $[0.0, 1.0]$ to define the following incremental utility function:

$$u^+(e, x'_j) = \frac{r(e, x'_j) + 1}{2}. \quad (7)$$

This can be applicable to both the direct utility functions: u_{OP}^+ and u_{DM}^+ .

4 Experimental Settings

This section describes our experimental settings; more comprehensive descriptions are in Appendix.

4.1 LLM and Prompt

We use Flan-PaLM 2 (L) (Google et al., 2023) as our LLM. This is an instruction-tuned model, and we follow Gao et al. (2023) to design the prompt $\text{prompt}(x, \mathcal{E}_x)$ that concatenates a task instruction, demonstrations \mathcal{E}_x , and an input x .

4.2 Tasks and Reward Functions

We focus on NLP datasets that are *not* used in the Flan instruction tuning (Chung et al., 2022), across different tasks and languages; Arabic (ar), English (en), Finnish (fi), Japanese (ja), and Spanish (es) are covered. Table 2 shows the dataset statistics, and Table 12 in Appendix shows some examples.

Binary classification The task is to output a single class label for binary detection, and the value of the reward function $R(y^*, y'_j)$ is 1.0 if $y^* = y'_j$, and 0.0 otherwise. The evaluation metric is corpus-level detection F1.

- **ISD** is a dataset for detection of “sarcasm” text in English and Arabic (Abu Farha et al., 2022).
- **EDOS-A** is a dataset for detection of “sexist” text in English (Kirk et al., 2023).

Multi-class classification The task is to output a single class label for fine-grained classification, and $R(y^*, y'_j)$ is the same as that of binary classification.

- **EDOS-B** is a dataset for 4-way fine-grained classification about the sexist text in English (Kirk et al., 2023); the evaluation metric is the macro F1.
- **CLINC** is a dataset for 150-way intent classification of user input text in English, which also tests out-of-domain detection (Larson et al., 2019); the evaluation metric is the joint accuracy proposed in Zhang et al. (2020).

Segmentation The task is to output a tagged version of an input text. $R(y^*, y'_j)$ is a word-level F1, and the evaluation metric is the same.

- **SSENT** is a decomposed subtask of structured sentiment analysis (Barnes et al., 2022), and we use the OpeNER portion (Agerri et al., 2013) in English and Spanish.

Translation The task is to translate text from a language to another, and $R(y^*, y'_j)$ is example-level GLEU (Wu et al., 2016). The evaluation metric is BLEU (Papineni et al., 2002).

- **XML-MT** is a dataset for translation of XML-tagged Web documentation (Hashimoto et al., 2019), and we take the English-to-Japanese/Finnish subsets.

4.3 Demonstration Retrieval and Reranking

As the off-the-shelf text retriever, we use a generic t5x retriever (Ni et al., 2022) used in previous work (Chaudhary et al., 2023; Gao et al., 2023), which has a multilingual capability based on mT5 (Xue et al., 2021). We also use mT5 (following RankT5 (Zhuang et al., 2023)) to train the

reranking model, which can be seamlessly applied to the different languages’ data.

We use $n = 10$ to train and validate the reranking model, and use $n = 50$ for the final test evaluation to investigate the generalization ability. We have run training of the reranking model in about 100 configurations in total, for different datasets, hyperparameter search of ℓ in Equation (6), and checkpoint selection with evaluation on the validation sets, before touching the test sets.

4.4 Methods to be Compared

For the evaluation, we report results by the following methods including the standard baselines and our reranking methods:

- **“0-shot”** is a baseline to know how the LLM performs only with the task instruction.
- **“RETR”** is another baseline to know how the LLM performs by simply selecting the top- k retrieved demonstrations.
- **“ u_{OP} ”** and **“ u_{DM} ”** are our main baselines by the reranking models trained with the direct utility functions (Section 3.1).
- **“ u_{OP}^+ ”** and **“ u_{DM}^+ ”** are the main methods for our analysis, by the reranking models trained with the incremental utility functions (Section 3.2).

When constructing the prompts, we order the demonstrations according to the retrieval scores for RETR, and according to the reranking scores for the reranking-based methods.

5 Results and Discussions

Table 3 shows the k -shot ICL results on the test sets, where we use $k = 1, 3, 5$ for all the datasets, except for XML-MT with $k = 1, 2, 3$ due to its prompt being too long. All the evaluation scores range in $[0, 100]$. We also conduct a cross-lingual experiment with SSENT; we use the English training examples for retrieval, and directly apply the English reranking models to the Spanish test set.

We first summarize general observations, and then show our analysis to provide insights.

- RETR ICL outperforms the 0-shot baseline on all the datasets, except for ISD (en). This indicates that the general input-text similarity would not always find useful demonstrations.

	Binary classification			Multi-class classification		
	ISD (en)	ISD (ar)	EDOS-A (en)	EDOS-B (en)	CLINC (en)	
0-shot	58.54	43.43	57.07	41.49	86.52	
RETR	55.35, 56.98, 58.61	53.07 , 55.81, 55.57	61.78, 65.74, 67.70	43.10, 46.14, 48.45	91.89, 92.12, 92.48	
u_{OP}^-	58.38 , 59.51, 59.84	52.92, 57.14 , 57.53	64.93, 69.06, 70.77	42.05, 44.58, 47.56	93.32, 94.30 , 94.22	
u_{OP}^+	58.11, 59.75 , 60.41	52.44, 55.87, 56.82	65.15 , 70.17 , 71.58	47.85 , 55.04 , 57.66	92.63, 93.20, 93.89	
u_{DM}^-	57.82, 58.97, 60.06	51.91, 55.61, 56.22	62.53, 66.07, 67.54	42.33, 45.48, 46.04	93.62 , 93.85, 93.55	
u_{DM}^+	57.36, 59.27, 59.09	51.73, 55.82, 57.09	63.75, 67.16, 69.26	42.89, 48.73, 51.12	92.01, 92.35, 92.68	
	Segmentation			Translation		
	SSENT (en)	SSENT (es)	SSENT (en→es)	XML-MT (ja)	XML-MT (fi)	
0-shot	34.06	27.88	27.88	48.04	33.45	
RETR	46.77, 50.60, 53.82	43.30, 49.71, 52.16	35.09, 36.59, 39.39	61.11, 62.80, 64.26	45.73, 47.50, 48.24	
u_{OP}^-	49.61, 55.42, 58.74	44.16, 51.89, 54.49	35.46, 36.73, 41.13	63.70, 65.35, 66.32	46.82, 48.10, 48.70	
u_{OP}^+	52.03 , 56.69 , 60.34	44.44, 53.56, 55.78	36.70, 39.58 , 42.27	63.53, 65.61, 66.14	47.67, 49.41 , 49.82	
u_{DM}^-	50.76, 53.53, 57.18	46.92 , 53.66 , 53.88	37.07 , 38.00, 40.49	64.20 , 66.09 , 66.49	48.01, 49.20, 49.88	
u_{DM}^+	49.15, 54.20, 57.74	45.54, 52.46, 53.43	35.54, 36.42, 40.14	64.02, 65.86, 66.56	48.34 , 49.26, 49.92	

Table 3: Test results. Each cell reports scores with $k = 1, 3, 5$ (except for XML-MT with $k = 1, 2, 3$).




Figure 2: Distributions of u_{DM} and u_{OP} on the training sets of the three classification datasets.




Figure 3: Distributions of u_{DM}^+ and u_{OP}^+ on the training sets of the three classification datasets.

- Reranking (with at least one utility function) improves upon RETR. Only one exception is ISD (ar) with $k = 1$, but $k = 3$ and $k = 5$ still show the benefit.
- The OP utility tends to be more effective than the DM utility, especially on the classification tasks, while we can see potential advantages of DM on SSENT (es) and XML-MT.
- Our proposed incremental utility tends to improve upon the direct utility, especially with the OP utility.
- English-based reranking is well transferred to Spanish on SSENT. This is encouraging, because it is often the case that we only have English resources for training.
- As motivated and expected in Section 3.2, the incremental utility functions work the best with $0.0 < \ell < 1.0$, especially $\ell = 0.8$ (more details in Appendix E).

5.1 Analysis on Classification Tasks

A clear advantage of OP on the classification task is to provide denser training signals than DM. Figure 2 (best viewed in color) shows distributions of u_{DM} and u_{OP} on the training sets of ISD (en), EDOS-B, and CLINC; the OP-based utility values are well distributed across the value range. Figure 3 then shows distributions of u_{DM}^+ and u_{OP}^+ , and we can see the consistent trend.

Existence of negative incremental values A notable observation is that not all the demonstrations bring positive incremental values, as shown in Figure 3. Based on the definition of the incremental utility function (Equation (7)), a demonstration has a *negative* effect when its incremental utility value is less than 0.5. It is interesting to see that even the top-10 retrieved candidates can have negative impacts in ICL.

When is incremental utility effective? Now that we have seen the difference of the distribu-

	Contrastive examples	Improvement by u_{OP}^+
EDOS-B	88.6%	+8.79
ISD (en)	62.8%	+0.18
EDOS-A	56.2%	+0.71
CLINC	49.1%	-0.71
ISD (ar)	37.9%	-0.82

Table 4: Importance of the contrastive examples (with OP) for the classification tasks.

	$n' = 10$	$n' = 30$	$n' = 50$
Contrastive examples	39.5%	49.3%	53.5%

Table 5: The effects of n' on ISD (ar).

	Original	Constrained
u_{OP}	52.92, 57.14, 57.53	53.31, 58.25 , 59.17
u_{OP}^+	52.44, 55.87, 56.82	54.00 , 57.29, 57.10

Table 6: Comparison between original (Table 3) and constrained retrieval strategies on ISD (ar).

tions of the OP and DM values, a remaining question is why the incremental utility is not always effective. Specifically, Table 3 shows that u_{OP}^+ performs worse than u_{OP} on ISD (ar) and CLINC. To explain the results, we count the number of *contrastive* training examples that meet the following criterion:

a training example e is considered to be contrastive, if its corresponding n ($= 10$) demonstration candidates cover both the $u(e, x'_j) > u_0(x'_j)$ and $u(e, x'_j) < u_0(x'_j)$ cases.

We expect such a training example to contribute to effective contrast of the demonstrations. Table 4 shows the counting results, along with the improvement (averaged across $k = 1, 3, 5$) over u_{OP} by u_{OP}^+ in Table 3, and we can see a positive correlation as expected.

Can we increase contrastive examples? We then discuss how to increase the number of the contrastive examples. Inspired by constrained retrieval (Gao et al., 2023), we modify our simple retrieval process with the following steps:

1. retrieving top- N ($N > n$) demonstration candidates for a training example e ,
2. selecting n' candidates according to the retrieval scores, by ensuring that all the class labels' examples are equally selected,

3. computing the direct utility scores of all the n' candidates, and
4. selecting n ($=10$) candidates to make e a contrastive example if possible.

The step 4. is optional to control the number of candidates to be used for training the reranking model. For ISD (ar), we set $N = 300^2$ and $n' = 50$. Table 5 shows how the number of the contrastive examples increases as expected, and Table 6 shows the ISD (ar) test evaluation results. It is notable that the constrained retrieval lets u_{OP}^+ perform the best with $k = 1$.

Necessity of compositional utility We see in Table 6 that u_{OP} still performs better than u_{OP}^+ with $k = 3, 5$; however, this is less controllable in our work because the utility values are defined for each demonstration independently. Therefore, it is a crucial next step to investigate compositional effects of adding multiple demonstrations (Gupta et al., 2023; Ye et al., 2023).

General instruction In summary, we provide the following instruction for classification tasks:

- using u_{OP} if the number of the contrastive training examples is low (e.g., less than 50%), and otherwise u_{OP}^+ , and
- using the constrained retrieval if there are much less contrastive training examples (e.g., less than 40%).

5.2 Analysis on Non-Classification Tasks

Figure 4 and Figure 5 (best viewed in color) show the distributions of the utility values on SSENT (es) and XML-MT (ja) as in Section 5.1. Unlike the classification tasks, the direct utility values are well distributed by the task-oriented reward.

Is incremental utility effective with OP? One observation is that 60-80% of the u_{OP} values fall into the $[0.0, 0.05)$ bucket; this is presumably because the LLM is not finetuned with the task formats of SSENT and XML-MT, and generating long texts leads to lower probability values in general (Wu et al., 2016).³ This would hinder the effectiveness of u_{OP} on the non-classification tasks,

² N is a dataset-dependent (or retriever-dependent) hyperparameter, and we suggest to set the value to maximize the class label coverage of the retrieved candidates.

³Even if we would be able to re-scale the output probability values by the output sequence lengths, it is not always possible to have access to token-level information of the outputs. For example, some of the recent LLMs only provide APIs.




Figure 4: Distributions of u_{DM} and u_{OP} on the training sets of the two non-classification datasets.

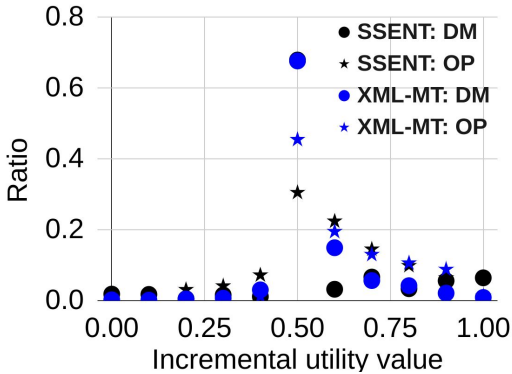


Figure 5: Distributions of u_{DM}^+ and u_{OP}^+ on the training sets of the two non-classification datasets.

and actually Table 3 shows that u_{DM} consistently outperforms u_{OP} (with $k = 1, 2, 3$) on XML-MT. However, as in the classification tasks, we can see that u_{OP}^+ helps improve the scores, and Table 7 shows the effects of the contrastive examples.

Limitation of DM We have seen the potential advantage of DM by the comparison between u_{DM} and u_{OP} on the SSENT and XML-MT datasets. This is an expected observation that u_{DM} works well on the non-classification tasks based on the nuanced values as the reward, only by having access to the output strings. However, u_{DM}^+ does not improve the scores in most of the cases (shown in Table 3). Table 8 shows the effects of the contrastive training examples from a viewpoint of DM. Surprisingly, there are much less contrastive examples than those with OP.

Necessity of compositional utility One reason why we could not effectively create the contrastive training examples is that, the LLM’s actual prediction is not easily affected by only one demonstration (i.e., $k = 1$). For example, if we inspect the LLM’s predictions on the validation set of SSENT

	Contrastive examples	Improvement by u_{OP}^+
SSENT (en)	56.4%	+1.76
SSENT (es)	56.4%	+1.08
XML-MT (fi)	44.0%	+1.09
XML-MT (ja)	30.2%	-0.03

Table 7: Importance of the contrastive examples (with OP) for the non-classification tasks.

	Contrastive examples	Improvement by u_{DM}^+
XML-MT (ja)	28.4%	-0.11
XML-MT (fi)	27.1%	+0.14
SSENT (es)	0.06%	-1.01
SSENT (en)	0.05%	-0.13

Table 8: Importance of the contrastive examples (with DM) for the non-classification tasks.

(en), only 27.4% of the independent demonstrations lead to different predictions than those with the 0-shot setting. If we include demonstrations that can modify the LLM’s predictions with $k = 2$, the ratio increases to 40.1%. We have then missed potentially useful signals for the utility estimation. This indicates that it is necessary for future work to estimate the utility of the demonstrations with $k > 1$ as also discussed in Section 5.1.

General instruction In summary, we provide the following instruction for non-classification tasks:

- using u_{DM} if well-defined nuanced metrics for the reward can be implemented, and
- considering the use of u_{OP}^+ otherwise, as in the classification tasks.

5.3 Demonstrations with High Utility

We have investigated the effectiveness of incremental utility, and in this section, we focus on how the utility values are related to the demonstrations’ ground-truth outputs. For the sake of simplicity, we take the classification datasets in this analysis.

For each dataset, we list all the $n|\mathcal{T}|$ input-demonstration pairs from the training set as in Section 2.2. We then sort the pairs to select the top- $m\%$ of them either by u_{OP} or u_{OP}^+ , and calculate the ratio of the pairs whose demonstration labels match with their paired input’s labels.

Table 9 shows the results, and overall, we can see that u_{OP} and u_{OP}^+ rank the pairs differently as expected in Section 3.2. Judging from Table 3

		Top-10%	Top-30%	Top-50%
ISD (en)	u_{OP}	68.7%	67.1%	64.8%
	u_{OP}^+	68.7%	67.2%	65.1%
ISD (ar)	u_{OP}	93.6%	89.4%	86.3%
	u_{OP}^+	81.1%	81.2%	81.7%
EDOS-A	u_{OP}	69.3%	66.7%	66.3%
	u_{OP}^+	72.5%	74.1%	73.7%
EDOS-B	u_{OP}	57.4%	53.8%	51.1%
	u_{OP}^+	93.3%	72.3%	62.8%
CLINC	u_{OP}	86.2%	86.9%	86.1%
	u_{OP}^+	85.3%	78.4%	78.7%

Table 9: Label matching ratio with high utility.

and Table 4, u_{OP}^+ performed the best on EDOS-B, and we can see much more matched labels when sorted by u_{OP}^+ than u_{OP} . This is considered to be a factor of the improvement, related to previous studies about the importance of the demonstrations’ labels (Min et al., 2022; Gao et al., 2023), while there are other factors like coverage (Gupta et al., 2023) or compositionality (Ye et al., 2023).

5.4 Generalization Ability

We have used Flan-PaLM 2 (L) as an LLM, and the t5x retriever as a baseline retriever; these are the core components in our evaluation pipeline. The last question we would like to answer here is whether our empirical findings are generalized by replacing the LLM or the retriever.

5.4.1 LLMs

We conduct experiments by simply replacing Flan-PaLM 2 (L) with another instruction-tuned model, Flan-T5 XXL (Chung et al., 2022), or a non-instruction-tuned model, GPT-J (Wang and Komatsuzaki, 2021). For Flan-T5, we use exactly the same prompt designs, and for GPT-J, we use different prompt designs as detailed in Appendix F.

Table 10 reports the results on EDOS-B, and we can see the consistent advantage of u_{OP}^+ on the dataset. One interesting observation is that GPT-J performs much worse than the instruction-tuned models in the baseline results, but the scores are dramatically improved by the reranked demonstrations. The results show that different LLMs make use of the demonstrations differently, which would be dependent on the target tasks, inherent biases of the LLMs, and the pre-training strategies. For example, the ‘‘Animosity’’ class on the dataset is much better predicted by GPT-J than the others. As a result, from a viewpoint of the macro F1 metric, GPT-J has an advantage, because performing poorly on a class significantly hurts the metric.

	Flan-T5 XXL	GPT-J
0-shot	35.28	19.78
RETR	36.59, 38.67, 38.73	37.12, 35.16, 29.68
u_{OP}	36.02, 38.81, 39.71	55.78, 57.23, 56.40
u_{OP}^+	41.30, 45.42, 46.89	58.97, 59.15, 58.54
u_{DM}	38.10, 40.13, 41.98	51.28, 53.08, 53.60
u_{DM}^+	36.85, 39.89, 40.70	52.33, 54.25, 54.55

Table 10: Results on EDOS-B by using Flan-T5 XXL or GPT-J as the LLM component.

	EDOS-B (en)	SSENT (en)
0-shot	41.49	34.06
RETR	40.81, 46.05, 49.47	45.20, 50.75, 54.43
u_{OP}	40.81, 43.57, 46.41	49.79, 54.75, 58.50
u_{OP}^+	44.70, 51.01, 53.31	51.60, 56.96, 60.21
u_{DM}	40.71, 44.00, 43.30	49.19, 54.15, 57.13
u_{DM}^+	42.04, 43.75, 46.18	48.14, 54.39, 57.95

Table 11: Results on EDOS-B and SSENT (en) by replacing the t5x retriever with a TF-IDF retriever only at the test time.

5.4.2 Retrievers

We then conduct another set of the experiments by replacing the t5x retriever with a TF-IDF retriever⁴ at the test time, to see how the reranking model (trained with the t5x retriever) works with different candidate sets. Table 11 shows the results on EDOS-B and SSENT (en). We can see that the reranking model works even with the TF-IDF-based candidates, showing the robustness of the reranking with the proposed utility functions.

6 Related Work

We discuss relationships between our work and previous work from several viewpoints. Interested readers may refer to a survey by Luo et al. (2024).

Utility function Recent work (Rubin et al., 2022; Li et al., 2023; Luo et al., 2023; Wang et al., 2023b) has proposed finetuning text retrieval models by estimating direct utility of a demonstration by u_{OP} -like utility functions. While they showed the effectiveness, we have made a deeper dive into understanding the effects of the different utility functions, especially with the enhance of incremental utility.

Loss function To finetune the retrieval models, some used contrastive learning to contrast a high-utility demonstration against a number of low-utility demonstrations (Rubin et al., 2022;

⁴https://scikit-learn.org/stable/modules/generated/sklearn.feature_extraction.text.TfidfVectorizer.html

Luo et al., 2023; Wang et al., 2023b), and others used ranking losses (Li et al., 2023). Contrastive learning (Karpukhin et al., 2020) and ranking losses (Zhuang et al., 2023) are widely used for document retrieval and reranking, but simple pointwise regression has shown to be effective as well (Nie et al., 2019; Nogueira and Cho, 2019; Asai et al., 2020). We employed the pointwise regression (Equation (2)), based on our observations that it significantly improves ICL upon the RETR baseline. Instead, we focused on the contrast between the 0-shot and 1-shot performances to derive the incremental utility functions.

Overall pipeline Lin et al. (2023) have taken one step further to finetune LLMs with a retriever, while our work has focused on using an LLM as it is. Within this framework, it is a common practice to directly finetune a dense text retrieval model (Karpukhin et al., 2020) for the demonstration retrieval (Rubin et al., 2022; Li et al., 2023; Luo et al., 2023), while Wang et al. (2023b) employed a two-step approach by knowledge distillation with a cross-attention reward model (similar to our reranking model). By contrast, we followed another major framework with retrieval and reranking (Zhuang et al., 2023), where we use a generic text retriever for all the different settings. This is useful in making our experiments controllable for analysis, in that the reranking models are always applied to the consistent candidate sets. One potential drawback of the cascaded pipeline is that the search space is limited; still, we have shown that there is large room for improvement even within the limited search space.

Constrained retrieval and selection This paper has focused on using feedback from LLMs to train the reranking models. Another line of related work is to improve the demonstration retrieval/selection process without finetuning additional models. Gao et al. (2023) take into account difficulty of each demonstration and entropy of the LLMs’ prediction, Mavromatis et al. (2023) select demonstrations based on uncertainty and diversity, Margatina et al. (2023) adopt active learning, and Gupta et al. (2023) propose to increase coverage of information about inputs. In Section 5.1, we have shown that a simple constrained retrieval has the potential to improve the reranking models, and it is an interesting direction to incorporate those kinds of techniques into the training of the reranking models.

Unified modeling We conducted our analysis on the different tasks in several languages, and all the reranking models were trained independently. Li et al. (2023) made an interesting attempt to finetune a retriever across different tasks, and showed promising results. Then another interesting future direction is to integrate our findings into such a unified framework; for example, instruction-based demonstration retriever/reranker is worth investigating by following Asai et al. (2023).

7 Conclusion

This paper has investigated how the output probability and downstream metric affect the utility estimation of demonstrations for ICL. Our in-depth analysis has shown that the output probability is robust on the classification tasks, and the downstream metric is robust especially when the output probability is not well distributed across the whole value range. Furthermore, we provided discussions about why and when our proposed incremental utility helps improve the ICL results.

Acknowledgements

First of all, we thank Aditi Chaudhary and Krishna Srinivasan for their contributions to the basics of the ICL work in our team. We also appreciate discussions with Lingyu Gao about her internship project and related work. Feedbacks from Zhuyun Dai and Tania Bedrax-Weiss were also valuable in polishing the draft. Lastly, we would like to thank anonymous ARR reviewers for their constructive feedbacks.

Limitations

Inference-time cost We adopted the retrieval-reranking framework, and the reranking model is a large cross-attention encoder model based on mT5 XXL. This is known to be much slower than another standard framework with dense retrievers (Karpukhin et al., 2020). The main goal in this paper is to investigate how to transfer the feedback signals from the LLMs to the demonstration selection, and this is still an open research question. Therefore, we first prioritized having an enough model capacity instead of restricting the inference-time cost, as in previous work on document retrieval (Asai et al., 2020). Once it gets practically useful, we can consider applying knowledge distillation (Wang et al., 2023b), or directly training dense retriever frameworks (Rubin et al., 2022).

Ethics Statement

Inherent biases It is known that LLMs have inherent biases that affect ICL results (Zhao et al., 2021; Min et al., 2022). We also observed that specific labels are more generated than others on the classification tasks, and the trends are different among different LLMs like PaLM 2 and GPT-J. It is then possible that the biases are transferred to the reranking models, and potentially we can further improve the results by debiasing the feedback signals. This has not been addressed in this paper, and is left for future work.

Contents of the datasets We used publicly available datasets that have been properly processed for privacy concerns. The EDOS dataset contains offensive contents as described in Kirk et al. (2023), and we intentionally avoided showing examples from the dataset. We do not intend to enhance the offensive contents in this paper, and instead solely focus on how the LLMs work on the challenging fine-grained classification.

References

- Ibrahim Abu Farha, Silviu Vlad Oprea, Steven Wilson, and Walid Magdy. 2022. [SemEval-2022 Task 6: iSarcasmEval, Intended Sarcasm Detection in English and Arabic](#). In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 802–814.
- Rodrigo Agerri, Montse Cuadros, Sean Gaines, and German Rigau. 2013. OpeNER: Open polarity enhanced named entity recognition. In *Sociedad Española para el Procesamiento del Lenguaje Natural*, volume 51, pages 215–218.
- Akari Asai, Kazuma Hashimoto, Hannaneh Hajishirzi, Richard Socher, and Caimeing Xiong. 2020. Learning to Retrieve Reasoning Paths over Wikipedia Graph for Question Answering. In *International Conference on Learning Representations*.
- Akari Asai, Timo Schick, Patrick Lewis, Xilun Chen, Gautier Izacard, Sebastian Riedel, Hannaneh Hajishirzi, and Wen-tau Yih. 2023. [Task-aware Retrieval with Instructions](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 3650–3675.
- Jeremy Barnes, Laura Ana Maria Oberländer, Enrica Troiano, Andrey Kutuzov, Jan Buchmann, Rodrigo Agerri, Lilja Øvrelid, and Erik Velldal. 2022. SemEval-2022 Task 10: Structured Sentiment Analysis. In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language Models are Few-Shot Learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Aditi Chaudhary, Karthik Raman, Krishna Srinivasan, Kazuma Hashimoto, Mike Bendersky, and Marc Najork. 2023. [Exploring the Viability of Synthetic Query Generation for Relevance Prediction](#). *arXiv preprint cs.IR 2305.11944*.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. [Scaling Instruction-Finetuned Language Models](#). *arXiv preprint cs.LG 2210.11416*.
- Alexandre de Brébisson and Pascal Vincent. 2016. The Z-loss: a shift and scale invariant classification loss belonging to the Spherical Family. *arXiv preprint arXiv:1604.08859*.
- Lingyu Gao, Aditi Chaudhary, Krishna Srinivasan, Kazuma Hashimoto, Karthik Raman, and Michael Bendersky. 2023. [Ambiguity-Aware In-Context Learning with Large Language Models](#). *arXiv preprint cs.CL 2309.07900*.
- Google, Rohan Anil, Andrew M. Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, Eric Chu, Jonathan H. Clark, Laurent El Shafey, Yanping Huang, Kathy Meier-Hellstern, Gaurav Mishra, Erica Moreira, Mark Omernick, Kevin Robinson, Sebastian Ruder, Yi Tay, Kefan Xiao, Yuanzhong Xu, Yujing Zhang, Gustavo Hernandez Abrego, Junwhan Ahn, Jacob Austin, Paul Barham, Jan Botha, James Bradbury, Siddhartha Brahma, Kevin Brooks, Michele Catasta, Yong Cheng, Colin Cherry, Christopher A. Choquette-Choo, Aakanksha Chowdhery, Clément Crepey, Shachi Dave, Mostafa Dehghani, Sunipa Dev, Jacob Devlin, Mark Diaz, Nan Du, Ethan Dyer, Vlad Feinberg, Fangxiaoyu Feng, Vlad Fienber, Markus Freitag, Xavier Garcia, Sebastian Gehrmann, Lucas Gonzalez, Guy Gur-Ari, Steven Hand, Hadi

- Hashemi, Le Hou, Joshua Howland, Andrea Hu, Jeffrey Hui, Jeremy Hurwitz, Michael Isard, Abe Ittycheriah, Matthew Jagielski, Wenhao Jia, Kathleen Kenealy, Maxim Krikun, Sneha Kudugunta, Chang Lan, Katherine Lee, Benjamin Lee, Eric Li, Music Li, Wei Li, YaGuang Li, Jian Li, Hyeontaek Lim, Hanzhao Lin, Zhongtao Liu, Frederick Liu, Marcello Maggioni, Aroma Mahendru, Joshua Maynez, Vedant Misra, Maysam Moussalem, Zachary Nado, John Nham, Eric Ni, Andrew Nystrom, Alicia Parrish, Marie Pellat, Martin Polacek, Alex Polozov, Reiner Pope, Siyuan Qiao, Emily Reif, Bryan Richter, Parker Riley, Alex Castro Ros, Aurko Roy, Brennan Saeta, Rajkumar Samuel, Renee Shelby, Ambrose Slone, Daniel Smilkov, David R. So, Daniel Sohn, Simon Tokumine, Dasha Valter, Vijay Vasudevan, Kiran Vodrahalli, Xuezhi Wang, Pidong Wang, Zirui Wang, Tao Wang, John Wieting, Yuhuai Wu, Kelvin Xu, Yunhan Xu, Linting Xue, Pengcheng Yin, Jiahui Yu, Qiao Zhang, Steven Zheng, Ce Zheng, Weikang Zhou, Denny Zhou, Slav Petrov, and Yonghui Wu. 2023. [PaLM 2 Technical Report](#).
- Shivanshu Gupta, Matt Gardner, and Sameer Singh. 2023. Coverage-based Example Selection for In-Context Learning. *arXiv preprint cs.CL 2305.14907*.
- Kazuma Hashimoto, Raffaella Buschiazzo, James Bradbury, Teresa Marshall, Richard Socher, and Caiming Xiong. 2019. [A High-Quality Multilingual Dataset for Structured Documentation Translation](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers)*, pages 116–127.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. [Survey of Hallucination in Natural Language Generation](#). *ACM Computing Surveys*, 55(12).
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense Passage Retrieval for Open-Domain Question Answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781.
- Hannah Kirk, Wenjie Yin, Bertie Vidgen, and Paul Röttger. 2023. [SemEval-2023 Task 10: Explainable Detection of Online Sexism](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, pages 2193–2210.
- Stefan Larson, Anish Mahendran, Joseph J. Peper, Christopher Clarke, Andrew Lee, Parker Hill, Jonathan K. Kummerfeld, Kevin Leach, Michael A. Laurenzano, Lingjia Tang, and Jason Mars. 2019. [An Evaluation Dataset for Intent Classification and Out-of-Scope Prediction](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1311–1316.
- Xiaonan Li, Kai Lv, Hang Yan, Tianyang Lin, Wei Zhu, Yuan Ni, Guotong Xie, Xiaoling Wang, and Xipeng Qiu. 2023. [Unified Demonstration Retriever for In-Context Learning](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4644–4668.
- Xiaonan Li and Xipeng Qiu. 2023. [Finding Support Examples for In-Context Learning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 6219–6235, Singapore. Association for Computational Linguistics.
- Xi Victoria Lin, Xilun Chen, Mingda Chen, Weijia Shi, Maria Lomeli, Rich James, Pedro Rodriguez, Jacob Kahn, Gergely Szilvassy, Mike Lewis, Luke Zettlemoyer, and Scott Yih. 2023. [RA-DIT: Retrieval-Augmented Dual Instruction Tuning](#). *arXiv preprint cs.CL 2310.01352*.
- Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2022. [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*, pages 100–114.
- Man Luo, Xin Xu, Zhuyun Dai, Panupong Pasupat, Mehran Kazemi, Chitta Baral, Vaiva Imbrasaitė, and Vincent Y Zhao. 2023. [Dr.ICL: Demonstration-Retrieved In-context Learning](#). *arXiv preprint cs.CL 2305.14128*.
- Man Luo, Xin Xu, Yue Liu, Panupong Pasupat, and Mehran Kazemi. 2024. [In-context Learning with Retrieved Demonstrations for Language Models: A Survey](#). *arXiv preprint arXiv:2401.11624*.
- Katerina Margatina, Timo Schick, Nikolaos Aletras, and Jane Dwivedi-Yu. 2023. [Active Learning Principles for In-Context Learning with Large Language Models](#). *arXiv preprint cs.CL 2305.14264*.
- Costas Mavromatis, Balasubramaniam Srinivasan, Zhengyuan Shen, Jian Zhang, Huzefa Rangwala, Christos Faloutsos, and George Karypis. 2023. [Which Examples to Annotate for In-Context Learning? Towards Effective and Efficient Selection](#). *arXiv preprint cs.CL 2310.20046*.
- Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. [Rethinking the Role of Demonstrations: What Makes In-Context Learning Work?](#) In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11048–11064.
- Jianmo Ni, Gustavo Hernandez Abrego, Noah Constant, Ji Ma, Keith Hall, Daniel Cer, and Yinfei

- Yang. 2022. [Sentence-T5: Scalable Sentence Encoders from Pre-trained Text-to-Text Models](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1864–1874.
- Yixin Nie, Songhe Wang, and Mohit Bansal. 2019. [Revealing the Importance of Semantic Retrieval for Machine Reading at Scale](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2553–2566.
- Rodrigo Nogueira and Kyunghyun Cho. 2019. [Passage Re-ranking with BERT](#). *arXiv preprint cs.IR 1901.04085*.
- OpenAI. 2023. [GPT-4 Technical Report](#). *arXiv preprint arXiv:2303.08774*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a Method for Automatic Evaluation of Machine Translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318.
- Karthik Raman, Iftekhar Naim, Jiecao Chen, Kazuma Hashimoto, Kiran Yalasang, and Krishna Srinivasan. 2022. [Transforming Sequence Tagging Into A Seq2Seq Task](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11856–11874.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2016. [Sequence Level Training with Recurrent Neural Networks](#). In *International Conference on Learning Representations*.
- Adam Roberts, Hyung Won Chung, Anselm Levskaya, Gaurav Mishra, James Bradbury, Daniel Andor, Sharan Narang, Brian Lester, Colin Gaffney, Afroz Mohiuddin, Curtis Hawthorne, Aitor Lewkowycz, Alex Salcianu, Marc van Zee, Jacob Austin, Sebastian Goodman, Livio Baldini Soares, Haitang Hu, Sasha Tsvyashchenko, Aakanksha Chowdhery, Jasmijn Bastings, Jannis Bulian, Xavier Garcia, Jianmo Ni, Andrew Chen, Kathleen Kenealy, Jonathan H. Clark, Stephan Lee, Dan Garrette, James Lee-Thorp, Colin Raffel, Noam Shazeer, Marvin Ritter, Maarten Bosma, Alexandre Passos, Jeremy Maitin-Shepard, Noah Fiedel, Mark Omernick, Brennan Saeta, Ryan Sepassi, Alexander Spiridonov, Joshua Newlan, and Andrea Gesmundo. 2022. [Scaling Up Models and Data with \$t5x\$ and `seqio`](#). *arXiv preprint arXiv:2203.17189*.
- Ohad Rubin, Jonathan Herzig, and Jonathan Berant. 2022. [Learning to retrieve prompts for in-context learning](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2655–2671.
- Noam Shazeer and Mitchell Stern. 2018. [Adafactor: Adaptive Learning Rates with Sublinear Memory Cost](#). In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 4596–4604. PMLR.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open Foundation and Fine-Tuned Chat Models](#).
- Ben Wang and Aran Komatsuzaki. 2021. [GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model](#). <https://github.com/kingoflolz/mesh-transformer-jax>.
- Lean Wang, Lei Li, Damai Dai, Deli Chen, Hao Zhou, Fandong Meng, Jie Zhou, and Xu Sun. 2023a. [Label Words are Anchors: An Information Flow Perspective for Understanding In-Context Learning](#). *arXiv preprint arXiv:2305.14160*.
- Liang Wang, Nan Yang, and Furu Wei. 2023b. [Learning to Retrieve In-Context Examples for Large Language Models](#). *arXiv preprint cs.CL 2307.07164*.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Łukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. [Google’s Neural Machine Translation System: Bridging the Gap between Human and Machine Translation](#). *arXiv preprint cs.CL 1609.08144*.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A Massively Multilingual Pre-trained Text-to-Text Transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498.

- Jiacheng Ye, Zhiyong Wu, Jiangtao Feng, Tao Yu, and Lingpeng Kong. 2023. Compositional Exemplars for In-Context Learning. In *Proceedings of the 40th International Conference on Machine Learning*, ICML'23. JMLR.org.
- Jianguo Zhang, Kazuma Hashimoto, Wenhao Liu, Chien-Sheng Wu, Yao Wan, Philip Yu, Richard Socher, and Caiming Xiong. 2020. Discriminative Nearest Neighbor Few-Shot Intent Detection by Transferring Natural Language Inference. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5064–5082.
- Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate before use: Improving few-shot performance of language models. In *International Conference on Machine Learning*, pages 12697–12706.
- Honglei Zhuang, Zhen Qin, Rolf Jagerman, Kai Hui, Ji Ma, Jing Lu, Jianmo Ni, Xuanhui Wang, and Michael Berdersky. 2023. RankT5: Fine-Tuning T5 for Text Ranking with Ranking Losses. In *SIGIR '23: Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2308–2313.

Appendix

A Prompt Design for Flan Models

We use the prompt template defined in [Gao et al. \(2023\)](#) for the instruction-tuned LLMs in our experiments. We then need to have a task definition for each dataset with an input-output format.

B Datasets

Table 12 shows examples of input-output pairs from the datasets used in our experiments.

B.1 ISD

The task definition of ISD is as follows:

```
The goal of this task is to
identify if an input text is
sarcastic or non-sarcastic.
```

Our initial analysis showed that the LLMs in our experiments have inherent bias towards generating “sarcastic” and “non-sarcastic” instead of “Sarcastic” and “Non-sarcastic,” and then we decided to use the all-lowercased label strings. The same task definition is shared in Arabic and English.

B.2 EDOS-A

The task definition of EDS-A is as follows:

```
The goal of this task is to
identify if an input text is
Sexist or Non-sexist.
```

In contrast to ISD, the LLMs in our experiments are biased towards generating “Sexist” and “Non-sexist” with the first letters capitalized.

B.3 EDOS-B

The task definition of EDOS-B is as follows:

```
The goal of this task is to
identify a category of a sexist
text from Threat, Derogation,
Animosity, or Prejudice.
```

The first letters are capitalized in the class labels with the same reason as in EDOS-A. [Gao et al. \(2023\)](#) incorporated more detailed descriptions of the fine-grained class labels into their task definition, while we simply listed the class labels to make all the tasks be tested under the consistent setting.

B.4 CLINC

The task definition of CLINC is as follows:

```
The goal of this task is to identify
a service domain and an intent
given a user input. There are
10 domains: "auto_and_commute"
"banking" "credit_cards" "home"
"kitchen_and_dining" "meta"
"small_talk" "travel" "utility"
"work" . For each domain, there are
15 intents: "auto_and_commute"=[
"current_location" "oil_change_when"
"oil_change_how" "uber" "traffic"
"tire_pressure" "schedule_maintenance"
"gas" "mpg" "distance" "directions"
"last_maintenance" "gas_type"
"tire_change" "jump_start" ],
"banking"=[ "freeze_account" "routing"
"pin_change" "bill_due" "pay_bill"
"account_blocked" "interest_rate"
"min_payment" "bill_balance"
"transfer" "order_checks" "balance"
"spending_history" "transactions"
"report_fraud" ], "credit_cards"=[
"replacement_card_duration"
"expiration_date" "damaged_card"
"improve_credit_score"
"report_lost_card" "card_declined"
"credit_limit_change" "apr"
"redeem_rewards" "credit_limit"
"rewards_balance" "application_status"
"credit_score" "new_card"
"international_fees" ], "home"=[
"what_song" "play_music"
"todo_list_update" "reminder"
"reminder_update" "calendar_update"
"order_status" "update_playlist"
"shopping_list" "calendar"
"next_song" "order" "todo_list"
"shopping_list_update" "smart_home"
], "kitchen_and_dining"=[
"food_last" "confirm_reservation"
"how_busy" "ingredients_list"
"calories" "nutrition_info"
"recipe" "restaurant_reviews"
"restaurant_reservation"
"meal_suggestion"
"restaurant_suggestion"
"cancel_reservation"
"ingredient_substitution"
```

Dataset	Input	Output
CLINC	what expression would i use to say i love you if i were an italian	travel: translate
SSENT	The food was very basic , but edible with good bread , soup and desserts .	The food was <Negative>very basic</Negative> , but <Positive>edible </Positive> with <Positive>good</Positive> bread , soup and desserts .
XML-MT	If you have permission to edit public templates, from Setup, enter <userinput>Email Templates</userinput> in the <parmname>Quick Find</parmname> box, then select <uicontrol>Classic Email Templates</uicontrol>.	Jos sinulla on oikeus muokata julkisia malleja, kirjoita Määrittäykset-valikon <parmname>Pikahaku</parmname>-kenttään <userinput>Sähköpostimallit </userinput> ja valitse <uicontrol>Classic-sähköpostimallit</uicontrol>.

Table 12: Examples of input-output pairs in the training sets.

```
"cook_time" "accept_reservations"
], "meta"=[ "change_speed"
"user_name" "whisper_mode" "yes"
"change_volume" "no" "change_language"
"repeat" "change_accent" "cancel"
"sync_device" "change_user_name"
"change_ai_name" "reset_settings"
"maybe" ], "small_talk"=[
"who_made_you" "meaning_of_life"
"who_do_you_work_for" "do_you_have_pets"
"what_are_your_hobbies" "fun_fact"
"what_is_your_name" "where_are_you_from"
"goodbye" "thank_you" "greeting"
"tell_joke" "are_you_a_bot"
"how_old_are_you" "what_can_i_ask_you"
], "travel"=[ "plug_type"
"travel_notification" "translate"
"flight_status" "international_visa"
"timezone" "exchange_rate"
"travel_suggestion" "travel_alert"
"vaccines" "lost_luggage" "book_flight"
"book_hotel" "carry_on" "car_rental"
], "utility"=[ "weather" "alarm" "date"
"find_phone" "share_location" "timer"
"make_call" "calculator" "definition"
"measurement_conversion" "flip_coin"
"spelling" "time" "roll_dice" "text"
], "work"=[ "pto_request_status"
"next_holiday" "insurance_change"
"insurance" "meeting_schedule" "payday"
"taxes" "income" "rollover_401k"
"pto_balance" "pto_request" "w2"
"schedule_meeting" "direct_deposit"
"pto_used" ]. Then the output is like
"domain: intent". If the input does
not belong to any of the domains, the
Answer is "oos".
```

The class labels are grouped according to the domains.⁵ The output label is with the domain name,

⁵<https://github.com/clinc/oos-eval/blob/master/data/domains.json>

but in the evaluation, we only refer to the intent class. The “oos” class is used to evaluate the detection of out-of-domain inputs (Larson et al., 2019).

In our preliminary experiments, we have observed that simply using the LLMs’ output strings leads very low recall of the “oos” class. This is mainly because the LLMs are biased towards generating in-domain class labels affected by in-domain demonstrations, which is consistent with Min et al. (2022) about LLMs’ bias in ICL. We then follow a threshold-based strategy in Zhang et al. (2020); we replace an output of an in-domain class label with “oos” if its prediction probability is below a pre-defined threshold. The threshold values are 0.6 for $k = 1$, 0.7 for $k = 2$, and 0.8 for $k = 3, 4, 5$.

B.5 SSENT

The task definition of SSENT is as follows:

The goal of this task is to copy the given hotel review by tagging sentiment-expressing phrases with the markup tags: <Positive></Positive> or <Negative></Negative>. Then the output is like "word1 <Positive>word2 word3</Positive> word4 <Negative>word5</Negative>". If there are not such tagged phrases, the Answer is Neutral.

For the output format, we followed Raman et al. (2022). The same task definition is shared in English and Spanish.

The structured sentiment analysis is a complex task in the original dataset (Barnes et al., 2022), where the task requires extracting tuples of (polar expression, subject, object) from an input text. In our preliminary experiments, we observed that, without finetuning, it is not trivial how to handle this complicated task in few-shot ICL.

We then decided to only focus on the polar ex-

pression extraction task as a first step, and the sub-task is considered to be a combination of sentiment classification and sequence labeling. It is an interesting future direction to investigate end-to-end evaluation of such a complex task in ICL.

B.6 XML-MT

The task definition of XML-MT is as follows:

```
The goal of this task is to
translate an XML-tagged text from
English to LANGUAGE by preservin
the XML structure. Both the
input and output are like "word1
<tag-A>word2 word3</tag-A> word4
<tag-B>word5</tag-B>".
```

The placeholder `LANGUAGE` is replaced with a language name (Finnish or Japanese in our experiments).

The test set has not been publicly released in the original dataset.⁶ We sampled 500 examples from the original validation set for our validation set, and used the rest (1,500 examples) for our test set.

C Hallucination in LLMs' Prediction

A concern about using the text generation models is hallucination (Ji et al., 2023). In general, we cannot perfectly prevent the models from generating texts with hallucinations. We do not discuss fact checking or semantics of the outputs, but instead solely focus on the output formats.

C.1 ISD

We did not observe any hallucinations on ISD in our experiments. This is because we carefully selected the label strings as discussed in Appendix B.1.

C.2 EDOS-A

We did not observe any hallucinations on EDOS-A in our experiments. This is also because we carefully selected the label strings as discussed in Appendix B.2.

C.3 EDOS-B

We observed a very small number of hallucinations on EDOS-B:

1. Sexual threat

⁶<https://github.com/salesforce/localization-xml-mt/tree/master/data>

2. Sexual Objectification

3. Objectification

The case 1. is simply mapped to `Threat`, and the cases 2. and 3. are mapped to `Derogation` according to the description of the label in Kirk et al. (2023). However, the frequency is low; for example, the cases 2. and 3. are observed with 2 out of 486 examples on the validation set.

Note that we never observed such hallucinations when we used Flan-T5 XXL in Section 5.4. Therefore, this can be different across different LLMs.

C.4 CLINC

We observed only one hallucination case on CLINC, out of the 3,100 validation examples. That case missed its domain label, only generating its intent label; in our evaluation pipeline, we automatically penalized such outputs that do not meet the defined output format. However, as mentioned above, there was only one such case in our experiments.

C.5 SSENT

For SSENT, we have tried several different input-output formats, before we decided to use the format described in the task definition. Specifically, the use of "Neutral" significantly reduced the chance to generate outputs with hallucinations. We could not perfectly mitigate hallucinations on SSET, though.

However, we observed that most of the hallucinations can be trivially fixed with simple rules, by referring to the original inputs. For example, the followings are typical cases:

- original: don 't → model output: do n't
- original: weren 't → model output: were n't
- original: wasn 't → model output: was n't

We can see that the model outputs look more natural than the original tokens, where the LLM unexpectedly *fixed* the unnatural tokenization. This is then more like a dataset issue.

Another type of hallucination was actually output formatting errors, such as missing a closing tag `</Positive>`. For such hallucinations that are not trivially fixed, our evaluation pipeline automatically replaces the outputs with "Neutral."

	$\ell = 0.0$	$\ell = 0.5$	$\ell = 0.8$	$\ell = 1.0$
u_{OP}^+	46.91	46.80	48.80	46.98

Table 13: The effects of changing the value of ℓ in Equation (6), on the EDOS-B validation set.

C.6 XML-MT

For XML-MT, only the potential hallucination is XML tag formatting issue as in SSENT. The formatting issue is rarely observed, but as in the original evaluation script,⁷ we penalize such predictions.

D Training of Reranking Model

We use the T5X code base (Roberts et al., 2022)⁸ for the reranking model. Among the available checkpoints, we use mT5 XXL⁹ for all the experiments. We finetune the original mT5 checkpoint to be used as a regression model (Zhuang et al., 2023), once we create a set of training examples described in Section 2.2.

We use the Adafactor optimizer (Shazeer and Stern, 2018), along with Z-loss regularization (de Brébisson and Vincent, 2016), a constant learning rate of 0.001, and a batch size of 256. We evaluate checkpoints after every 100 updates of the model parameters; for the checkpoint selection, we use the reranking model to select demonstrations on the validation set of each dataset, and then select a checkpoint that leads to the best downstream metric with the LLMs’ prediction.

To save the inference time of the LLMs on the validation sets, we pre-compute and cache all the possible k -shot predictions beforehand. For each validation example, the number of the demonstration combinations is ${}_n C_k$, where demonstrations are ordered according to the original retrieval scores. This caching strategy allows us to quickly check the LLMs’ performance without repeating the inference multiple times.

E Effects of ℓ in Incremental Utility

In the beginning of Section 5, we have mentioned that $\ell = 0.8$ in Equation (6) is recommended to implement the incremental utility function. We report

⁷<https://github.com/salesforce/localization-xml-mt/blob/master/scripts/evaluate.py>

⁸<https://github.com/google-research/t5x>

⁹<https://github.com/google-research/t5x/blob/main/docs/models.md#mt5-checkpoints>

	$\ell = 0.0$	$\ell = 0.5$	$\ell = 0.8$	$\ell = 1.0$
u_{OP}^+	50.22	48.13	50.83	45.40
u_{DM}^+	49.11	49.38	49.82	45.18

Table 14: The effects of changing the value of ℓ in Equation (6), on the SSENT (es) validation set.

results of tuning the value of ℓ on a classification dataset, EDOS-B, and a non-classification dataset, SSENT (es).

Table 13 and Table 14 show the results on the validation sets. Note that changing the values of ℓ does not affect the results with u_{DM}^+ for the classification tasks, because the original reward values are binary (i.e., 0.0 or 1.0). We can see that $\ell = 0.8$ leads to the best scores, presumably because it takes a good balance between $\ell = 0.0$ and $\ell = 1.0$ as motivated in Section 3.2.

F Prompt Design for GPT-J

GPT-J is a publicly available variant of the GPT models,¹⁰ and it has been used in previous work (Wang et al., 2023a). GPT-J is not an instruction-tuned model, and we design our prompt in a question-answering format:

Text: `INPUT TEXT\n`
Question: `which sexist category does the text express, Animosity, Derogation, Prejudice, or Threat?\n`
Answer: `the sexist category is`

The placeholder INPUT TEXT is replaced with an actual input text, and this corresponds to the 0-shot inference. Then, a prompt in 2-shot ICL is as follows:

Text: `DEMONSTRATION INPUT 1\n`
Question: `which sexist category does the text express, Animosity, Derogation, Prejudice, or Threat?\n`
Answer: `the sexist category is DEMONSTRATION OUTPUT 1\n`
Text: `DEMONSTRATION INPUT 2\n`
Question: `which sexist category does the text express, Animosity, Derogation, Prejudice, or`

¹⁰We used the transformers library for GPT-J: https://huggingface.co/docs/transformers/model_doc/gptj.

Threat?\n

Answer: the sexist category is
DEMONSTRATION OUTPUT 2\n

Text: INPUT TEXT\n

Question: which sexist category
does the text express, Animosity,
Derogation, Prejudice, or
Threat?\n

Answer: the sexist category is

Once the prompt is designed, we can seamlessly
run the experiments with GPT-J.

G Compute Resources

Flan-PaLM 2 (L) Readers may refer to [Google et al. \(2023\)](#).

Flan-T5 XXL The size of Flan-T5 XXL is 11B ([Chung et al., 2022](#)), and we used 64 v3 TPU chips to run inference with it.

GPT-J The size of GPT-J is 6B ([Wang and Komatsuzaki, 2021](#)), and we used one NVIDIA A100 GPU to run inference with it.

mT5 XXL for reranking The size of mT5 XXL is 13B ([Xue et al., 2021](#)), and we used 64 v3 TPU chips to run both training and inference with it for reranking.