

The Alchemy of Thought: Understanding In-Context Learning Through Supervised Classification

Harshita Narnoli and Mihai Surdeanu

Department of Computer Science, University of Arizona, Tucson, AZ, USA
{harshitanarnoli, msurdeanu}@arizona.edu

Abstract

In-context learning (ICL) has become a prominent paradigm to rapidly customize LLMs to new tasks without fine-tuning. However, despite the empirical evidence of its usefulness, we still do not truly understand how ICL works. In this paper, we compare the behavior of in-context learning with supervised classifiers trained on ICL demonstrations to investigate three research questions: (1) Do LLMs with ICL behave similarly to classifiers trained on the same examples? (2) If so, which classifiers are closer, those based on gradient descent (GD) or those based on k-nearest neighbors (kNN)? (3) When they do not behave similarly, what conditions are associated with differences in behavior? Using text classification as a use case, with six datasets and three LLMs, we observe that LLMs behave similarly to these classifiers when the relevance of demonstrations is high. On average, ICL is closer to kNN than logistic regression, giving empirical evidence that the attention mechanism behaves more similarly to kNN than GD. However, when demonstration relevance is low, LLMs perform better than these classifiers, likely because LLMs can back off to their parametric memory, a luxury these classifiers do not have.

1 Introduction

In-context learning (ICL) has emerged as the dominant paradigm for adapting large language models (LLMs) to new tasks without requiring any parameter updates. ICL achieves this by including a few relevant examples (or demonstrations) in the LLM’s prompt (Brown et al., 2020; Wei et al., 2022). However, despite the considerable amount of empirical evidence for ICL’s usefulness, the reasons behind its success remain unclear. On one hand, a few notable works establish an equivalence between ICL and gradient descent (GD) based on synthetic settings (Von Oswald et al., 2023; von Oswald et al., 2024; Akyürek et al., 2023; Ahn et al.,




Figure 1: Performance on six datasets of kNN, LR, GPT-4o-mini, Llama-3.1-8B-Instruct, and Qwen-7B-Chat models for the same $k = 10$ examples used by each model. In several datasets, LLMs perform similarly to kNN and LR, but this is not always the case.

2023). These works hypothesize that transformers use internal optimization using gradient-based algorithms to learn in-context. On the other hand, Deutch et al. (2024) observes a weak correlation between ICL and GD, uncovering several major discrepancies in the flow of information throughout the model between the two processes. These contradictions suggest that we still do not yet know *how* ICL works.

In this paper, we continue the quest to understand ICL’s behavior in the context of text classification tasks. Unlike previous work, we base our work on the intuitive empirical observation that ICL *should* be similar to supervised classifiers trained on ICL examples.¹ We further refine this idea into two directions. First, building on the above hypothesis that ICL may behave similar to GD, these classifiers could belong to the family of classifiers based on gradient descent. To explore this direction, we leverage Logistic Regression (LR) (LaValley, 2008). Second, we observe that the decoder’s attention mechanism feels sim-

¹For brevity, we refer to these classifiers trained on ICL examples simply as *classifiers* from here on. We use *example* or *demonstration* interchangeably to point to the demonstrations used in ICL or the training examples used by the classifiers.

ilar to the example weights assigned by weighted k-nearest-neighbors (kNN) (Dudani, 1976). Based on this, we explore two variants of kNN classifiers. We convert this intuition into three research questions: (1) *Do LLMs with ICL behave similarly to classifiers trained on the same examples?* (2) *If so, which classifiers are closer, those based on gradient descent or those based on kNN?* (3) *When they do not behave similarly, what conditions make them behave differently?*

To answer the above questions, we compare the performance and behavior of three LLMs with ICL (GPT-4o-mini (OpenAI, 2024), Llama-3.1-8B-Instruct (AI@Meta, 2024), and Qwen-7B-Chat (Yang et al., 2024)) against kNN and LR on six datasets. For a fair comparison, all models have access to the same training examples/demonstrations for each prediction. These examples were selected based on cosine similarity with the corresponding test example, using contextualized embeddings generated by Sentence-BERT (Reimers and Gurevych, 2019). In this paper, we focus on text-classification tasks due to their ubiquity and ease of manual analysis.

We draw several observations from our analysis. First, in most situations, LLMs obtain similar or close performance to the classifiers (see the three left-most datasets in Figure 1). More interestingly, LLMs *behave* similarly to kNN and LR in these situations, as supported by a high Kappa agreement (Cohen, 1960) between their outputs.

Second, LLMs show a slightly higher Kappa agreement with kNN, giving more empirical weight to the hypothesis that ICL behaves more similarly to the kNN mechanism.²

Third, these observations do not always hold (see the three right-most datasets in Figure 1). Our hypothesis is that the similarity in behavior is driven by the relevance of ICL examples, where *example relevance* characterizes the extent to which ICL examples align with the testing data point.³ To verify the hypothesis, we propose the following experiment: (a) We first annotate examples based on their relevance to the corresponding data point. The annotation is performed both manually and automatically. (b) Using this relevance, we determine whether its values are associated with the similarity between classifiers and LLMs. To do this, we

²However, please refer to Limitations for an additional discussion.

³For simplicity, henceforth we refer to example relevance as ‘relevance’.

compute the Kappa agreement between the classifier and LLM outputs. (c) We then examine the correlation between the Kappa values and the relevance scores. As anticipated, we observe a high correlation between relevance (both manual and machine) and Kappa agreement between models. In other words, when the relevance score is high, it naturally aligns with a higher Kappa agreement. This suggests that LLMs with ICL behave similarly to classifiers when the underlying examples are of high relevance. When not, their performance and behavior begin to diverge.

Beyond the more theoretical understanding of the behavior of ICL, our analysis yields an important practical insight: when relevant examples are available for a given task, LLMs can be replaced by a simpler and much more efficient classifier.

2 Related Work

We organize our discussion of previous work into two categories: papers that aim to understand how ICL works and those that improve its performance.

2.1 Understanding In-context Learning

Several research studies have focused on ways to understand ICL in different scenarios. In the context of gradient descent (GD), Dai et al. (2023) interpreted ICL as meta-optimizers through comparative analysis between ICL and explicit fine-tuning. Akyürek et al. (2023) theoretically proved that transformers can implement learning algorithms for linear models using gradient descent and closed-form ridge regression.

Around the same time, Von Oswald et al. (2023) explored the relation between training transformers on auto-regressive objectives and gradient-based meta-learning formulations. They investigated how transformers construct a loss function from given examples and the mechanisms through which they incorporate knowledge using the gradients of this loss function. Building on this idea, von Oswald et al. (2024) explored the underlying mechanisms of gradient-based mesa-optimization, i.e., a form of learning that “is acquired through training, as opposed to being inherent to the model” in autoregressive transformers trained for sequence modeling tasks. By reverse-engineering these models, they demonstrated that their in-context learning (ICL) capability arises from the mesa-layer, an innovative attention layer designed to efficiently solve a least-squares optimization problem. In the same space,

Zhou et al. (2025) analyzes whether LLMs perform ICL through an error-driven process similar to how humans learn. By testing for the inverse frequency effect, a phenomenon where rare examples more influence an agent than likely ones, the authors show that LLMs exhibit similar behavior to humans, suggesting that ICL involves an implicit error-driven learning mechanism like gradient descent.

In contrast, Deutch et al. (2024) reexamined the hypothesis that GD approximates ICL, identifying key issues in the evaluation metrics and baselines proposed by Dai et al. (2023). Shen et al. (2024) similarly points out the inconsistent behavior of ICL and GD in a natural setting. They demonstrated that the hand-constructed weights used in these studies show properties that differ from those found in practical scenarios. Fu et al. (2024) provided evidence that transformers acquire in-context learning (ICL) capabilities by employing higher-order optimization instead of gradient descent (GD). They demonstrated that transformer architectures can efficiently implement Newton’s method, in which transformers significantly outperformed GD in terms of speed, exhibiting exponential acceleration. Nafar et al. (2025) investigates the underlying mechanisms of ICL in LLMs within regression tasks. The work explores how LLMs learn from examples given in prompts versus relying on their internal knowledge.

Similar to most of these efforts, our work focuses on understanding ICL. We empirically observe that ICL behaves similarly to both gradient descent and kNN classifiers trained on the same examples as the demonstrations provided in ICL prompts. However, this observation only holds when the relevance of demonstrations is high. We believe this is more intuitive and has more practical implications.

2.2 Improving In-context Learning

Another line of research, such as Wang et al. (2024), examines the in-context learning phenomenon from a Bayesian perspective, viewing real-world LLMs as latent variable models. Based on this concept, the work proposes a practical and better demonstration selection algorithm using a small LLM to approximate optimal latent concepts, which are then generalized to larger LLMs. The method yields significant performance gains, bridging the gap between theory and practical LLM behavior. Focusing on the quality of examples, An et al. (2023) proposes SKILL-KNN, which is a training-free, skill-based approach for few-shot example selection in

in-context learning. The work utilizes prompting LLMs to generate task-relevant skill descriptions from raw inputs. SKILL-KNN mitigates biases associated with surface-level features commonly found in embedding-based selection methods to generate better examples. Whereas this prior research primarily uses kNN to find better examples for ICL, our work treats kNN as a comparative framework to gain insights into understanding the behavior of ICL.

Another work proposes framing the selection of demonstration examples as a sequential decision-making problem and employs reinforcement learning (RL) to learn policies for selecting examples (Zhang et al., 2022). This method is runtime expensive and does not consider examples based on similarity to the test data point. Instead, they observe that label distribution in the demonstration pool matters. In contrast, our work is simple, i.e., we use a standard model, Sentence-BERT, to provide a range of quality examples based on similarity. Further and more importantly, our work focuses on how the relevance of the examples impacts the decision of the downstream models.

Min et al. (2022) focuses on the structure of demonstrations rather than demonstration relevance for the test data. This work shows that factors like formatting and input-label alignment significantly affect LLM performance. Nevertheless, this approach is different than our work, as our main focus is to investigate the importance of demonstration relevance.

Additionally, Margatina et al. (2023) shows that selecting the semantically similar examples for prompts—using active learning strategies—can enhance ICL performance more than random demonstrations. It shows that the semantically similar demonstrations outperform other strategies and provide high-quality examples. Our paper takes this idea a step further. We aim to understand what happens when we use Sentence-BERT to select examples. Our work is different because we observe two important and practical conclusions when we use such examples: (i) ICL behaves very closely to kNN, which is novel, in cases of relevant examples; (ii) if the relevance of examples is not high, we do not observe the above correlation. This suggests that LLMs may back off to their existing knowledge in this situation.

3 Empirical Analysis

Despite the progress in understanding the behavior of ICL, a critical gap remains in how these systems behave the way they do. We design our goal of understanding ICL’s behavior as a comparison to classifiers trained on the same examples as ICL demonstrations. We approach this by performing a formal empirical analysis of the two, aimed at analyzing their similarity in performance and outputs across multiple models and datasets. Our findings revealed that ICL and classifiers are similar in performance, and in most cases, they generate the same labels.

The observation sparked our curiosity: When they do not behave similarly, what conditions are correlated with differences in behavior? To explore this, we annotate a sample of ICL examples with respect to their relevance and compute the correlation between example relevance and behavior similarity. The results are striking; we observe that LLMs behave similarly to the classifiers when the relevance of the demonstrations is high. We further investigate how changes in the relevance of examples influence the similarity in behavior.

3.1 Preliminaries

This section outlines the experimental setup used to support our findings. To provide a comprehensive overview, we perform text classification tasks on six datasets with kNN, LR, and three LLMs (GPT, Llama, and Qwen). Prior to running the classification experiments, we generate text embeddings for both the training and test inputs of each dataset using Sentence-BERT (Reimers and Gurevych, 2019).

3.1.1 Models

Sentence-BERT: The model generates embeddings that are used to retrieve the closest k examples from the training set based on similarity scores for each test case. To achieve this, we generate and save the contextualized embeddings for both the training and test data of each dataset using Sentence-BERT. During evaluation, we select the top- k training examples with the highest similarity by sorting them in descending order of cosine similarity between their embedding vectors and the corresponding test embedding. Note that, to have a fair comparison between classifiers and LLMs with ICL, the same examples were used by all LLMs.⁴

⁴Our goal in this paper is not to develop a novel example selection strategy, but rather to understand the classifiers and

kNN: We perform our experiment with two variants of kNN: *unweighted kNN* (Fix and Hodges, 1951) and *weighted kNN* (Dudani, 1976). In the unweighted case, the predicted label for a test point is determined by a simple majority vote, meaning the label most common among its nearest neighbors is chosen. In weighted kNN, however, each neighbor’s vote is scaled by its similarity score, giving closer or more similar examples greater influence in the decision. In our experiments, we observed that both approaches behave quite similarly, with nearly identical performance outcomes. Hence, for simplicity, we report the results of correlation using only the unweighted kNN.⁵

LR: Similar to kNN, we train a *Logistic Regression* model using labeled training data with ICL demonstrations when $k > 1$. In the case of a single training example—a situation that this discriminative classifier cannot handle—the model is configured to behave as a 1-NN classifier to keep the experiments consistent with kNN. We consider LR, as it is inspired by prior work on probing models that employ lightweight classifiers (Hewitt and Liang, 2019; Zhao and Bethard, 2020).

LLMs: We used three pre-trained LLMs: GPT-4o-mini (OpenAI, 2024), Llama-3.1-8B-Instruct (AI@Meta, 2024), and Qwen-7B-Chat (Yang et al., 2024). These LLMs have demonstrated proficiency across a range of tasks, such as text classification and sentiment analysis, using only a limited set of task-specific examples. We perform text classification using a common prompt with small changes based on the behavior of the LLM (a prompt example is provided in Appendix D). The LLMs range in size from 7B to 8B. As mentioned, all LLMs used exactly the same ICL demonstrations as the other classifiers.

3.1.2 Datasets

For our experiments, we used six datasets: AG News (Zhang et al., 2015), News Category (Misra, 2022), AuTextification (Sarvazyan et al., 2023), Implicit Prompts (Das et al., 2024), SST-2 (Socher et al., 2013), and SemEval-2019 Task 6 (Zampieri et al., 2019). During evaluation, for datasets with large test partitions, we conduct experiments using

LLM behavior under a shared, common strategy. It is quite possible that both classifiers and LLM absolute performance would improve with a more sophisticated example selection strategy.

⁵The performance results of weighted kNN are provided in Appendix G.

Dataset	$k = 1$					$k = 10$					$k = 20$					$k = 30$					
	kNN	LR	GPT	Llama	Qwen	kNN	LR	GPT	Llama	Qwen	kNN	LR	GPT	Llama	Qwen	kNN	LR	GPT	Llama	Qwen	
AG News	0.85	0.85	0.83	0.84	0.85	0.88	0.88	0.89	0.88	0.88	0.88	0.88	0.89	0.89	0.88	0.89	0.87	0.89	0.89	0.89	0.88
News Category	0.80	0.80	0.86	0.81	0.73	0.82	0.83	0.87	0.86	0.82	0.82	0.82	0.84	0.87	0.87	0.82	0.82	0.83	0.88	0.88	0.82
AuTexTification	0.54	0.54	0.52	0.55	0.54	0.59	0.58	0.58	0.54	0.52	0.58	0.58	0.58	0.55	0.48	0.59	0.58	0.55	0.54	0.52	
Implicit Prompts	0.69	0.69	0.74	0.58	0.41	0.78	0.78	0.78	0.71	0.50	0.79	0.78	0.78	0.74	0.44	0.79	0.79	0.78	0.73	0.48	
SST-2	0.68	0.68	0.92	0.80	0.94	0.72	0.73	0.92	0.90	0.93	0.73	0.73	0.92	0.92	0.93	0.73	0.74	0.92	0.92	0.94	
SemEval-2019	0.39	0.39	0.58	0.51	0.49	0.46	0.56	0.70	0.62	0.59	0.48	0.61	0.73	0.65	0.57	0.47	0.63	0.74	0.66	0.56	

Table 1: Test accuracy on six datasets for unweighted-kNN, LR, and three LLMs, using 1, 10, 20, and 30 examples.

a sample of approximately 1,000 test data points. Table 6 in Appendix B shows the total number of test cases along with the selected sample size and the number of labels. Additionally, the labels for these datasets are listed in Appendix B in Table 7.

3.2 Comparing Model Behavior

We begin with the analysis that motivates our study, highlighting the first research question: **RQ1: Do LLMs with ICL behave similarly to classifiers trained on the same examples?** To explore this, we set out to compare the behavior of three different LLMs (GPT-4o-mini, Llama-3.1-8B-Instruct, and Qwen-7B-Chat) with classifiers (kNN and LR).

3.2.1 Measuring Performance in Terms of Accuracy

We conduct a systematic analysis of all the models to measure performance, which is evaluated on different datasets with the same k examples. On examining their performance from Table 1, we see that the accuracy of kNN, LR, and three LLMs is similar in most cases. This is further visualized in a heatmap plot (Figure 2) that reflects the performance differences between kNN and each LLM, as well as between LR and each LLM, across all settings. The figure indicates that the accuracy differences between classifiers and the three LLMs are minuscule on three datasets (left part of the figure). For example, the performance difference on AG News and News Category is 0 or very close to 0 in most cases. However, these differences increase significantly on the three datasets in the right part of the figure, with the difference reaching as high as 35% between kNN and Qwen on Implicit Prompts.

Building on these findings, a natural next question arose; **RQ2: If ICL behaves similarly to classifiers with the same set of examples, which classifiers are closer, those based on gradient**

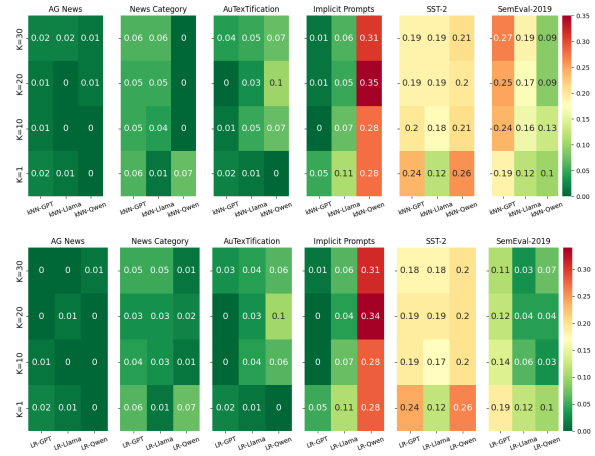


Figure 2: Visualization highlighting the difference in accuracy between kNN (top) and LR (bottom), each compared with three LLMs. Each vertical block is a different dataset; each row is a setting with a different number of ICL examples (from 1 to 30). The difference is also encoded with a color gradient, where green means lower difference (higher similarity) and red indicates high differences (low similarity).

descent (GD) or those based on k-nearest neighbors (kNN)? We observe that kNN and Logistic Regression have comparable performance, while ICL aligns more closely with kNN.⁶ However, as we have a very simple gradient-based classifier, the results may vary with non-linear GD classifiers.

3.2.2 Assessing Model Similarity Using Contingency Matrices

At this point, one might reasonably argue that performance parity among models does not necessarily imply similar behavior, as they may achieve comparable performance scores while exhibiting different predictions of labels. To better understand the similarity in performance, we next construct contingency matrices between each LLM

⁶Detailed comparison of kNN with LR is given in Appendix H.

and the classifiers' outputs. This helps to investigate whether classifiers and LLMs actually *behave* in a similar manner, i.e., *do classifiers and LLMs generate identical labels given the same input?*

In our work, we plot the contingency matrices between kNN and LR (X-axis) and the three LLMs (Y-axis) on the six datasets. These matrices provide a structured way to compare and evaluate the labels predicted by different models. That is, the more similar two outputs are, the more weight the diagonal will have. The matrices support our observation: classifiers and LLMs exhibit similar behavior on some datasets but not on all. For example, referring to the contingency matrices of AG News (Figure 3) and News Category (Figure 4), we find that they follow diagonal dominance, which indicates that both classifiers produce almost the same labels as LLMs. In an opposing case of Implicit Prompts (Figure 6), we observe that the classifiers tend to favor the offensive label. We hypothesize that this occurs because the label distribution of the dataset is heavily biased towards offensive prompts.⁷

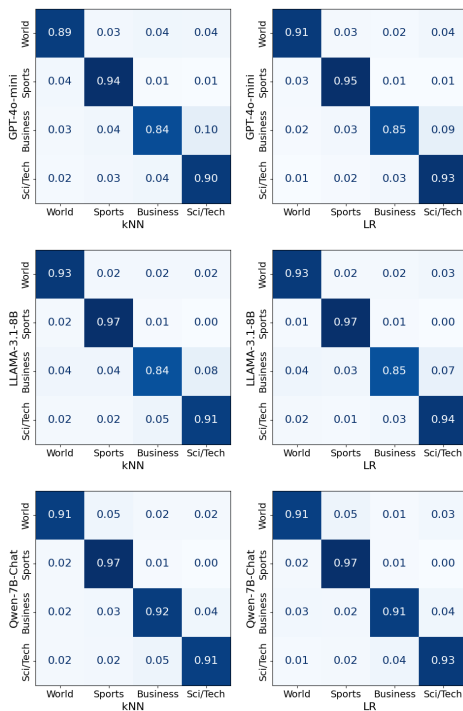


Figure 3: Contingency matrices for AG News: Comparison of kNN vs GPT (top-left), kNN vs Llama (middle-left), and kNN vs Qwen (bottom-left) and LR vs GPT (top-right), LR vs Llama (middle-right), and LR vs Qwen (bottom-right) for $k = 20$.

Overall, we observe that diagonal dominance

⁷All contingency matrices and extended discussion included in Appendix F.

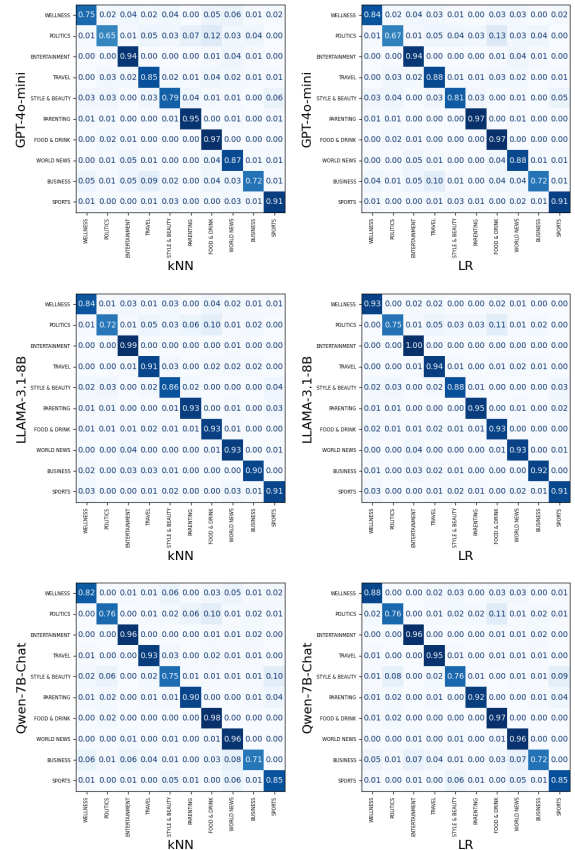


Figure 4: Contingency matrices for News Category: The same configuration follows from Figure 3.

is excellent in 2 out of 6 datasets, good in 2, and decent in 2.

3.3 Does Example Relevance Correlate with Behavior Similarity?

All the previous observations immediately indicate a puzzling conclusion: classifiers and LLMs with ICL obtain similar performance in general but not in all cases. These results lead to a new question: **RQ3: When LLMs and classifiers do not behave similarly, what conditions are correlated with differences in behavior?** To investigate this, we hypothesize that, given the performance of kNN is driven by examples, it is essential to analyze the correlation between example relevance and the performance of models.

3.3.1 Understanding Example Relevance

In our context, we define *example relevance* in terms of *task relevance for the provided evaluation data point*. Informally, relevance indicates how well a semantically similar example closely matches the input test data prompt for the task at hand. By selecting examples that are relevant to

the input prompt, we can measure whether a direct correlation between example relevance and the similarity in behavior of models exists.

To establish correlation, we mark the top- k semantically similar examples with a Boolean label for each example. We assign a value of ‘1’ to the examples that are relevant to the input test data prompt, and ‘0’ to those that are not. For each test data prompt, we determine an average of these values for all examples, which we call a *relevance score*, reflecting the percentage of the top- k examples that are relevant to the input prompt.

For labeling the examples, we consider two options: one based on manual annotations and the other based on the respective LLM. In the first case, considering human annotators for all examples across the entire test prompts is both challenging and expensive. For this reason, we sample and annotate a set of 50 test data prompts for each setting of $k = 1, 10,$ and 20 examples manually by Human-A.⁸ To verify the robustness of the manual annotations, a smaller subsample of data points was independently annotated by another human annotator (Human-B).⁹ We find substantial Kappa agreement ($\approx 70\%$) of relevance scores between the two human annotators.

We use the same set of 1,550 examples, this time using the respective LLM to calculate their corresponding relevance score. We then compare the agreement between relevance scores by human and machine annotators. To measure agreement, we introduce two scores. The first score, $\mathcal{K}_{H,M}^{rel}$ represents the Kappa agreement on relevance between a human and an LLM. The second measures the inclusion score (\mathcal{I}_{HCM}), which reflects the fraction of humans’ relevance scores that is included within that of LLMs. \mathcal{I}_{HCM} score of 0 indicates no inclusion, and a score of 1 indicates complete inclusion. Table 2 shows the Kappa ($\mathcal{K}_{H,M}^{rel}$), and Inclusion (\mathcal{I}_{HCM}) scores of relevance between Human-A and Gpt-4o-mini.

We find that even though the $\mathcal{K}_{H,M}^{rel}$ value is not good, the \mathcal{I}_{HCM} reflects an interesting conclusion that a vast majority of examples annotated as relevant by the human annotator were also considered relevant by the LLM annotator. In other words, the human annotators are more conservative than

⁸We annotated a total of $50 \times (20 + 10 + 1) = 1,550$ examples per dataset.

⁹The two annotators each annotated two data points per configuration per dataset, for a total of $2 \times (20 + 10) = 60$ examples per dataset.

Datasets	$k = 1$		$k = 10$		$k = 20$	
	$\mathcal{K}_{H,M}^{rel}$	\mathcal{I}_{HCM}	$\mathcal{K}_{H,M}^{rel}$	\mathcal{I}_{HCM}	$\mathcal{K}_{H,M}^{rel}$	\mathcal{I}_{HCM}
AG News	0	1	0.05	0.98	0.09	0.97
News Category	0.41	0.95	0.28	0.92	0.24	0.89
AuTexTification	0.52	0.92	0.25	0.76	0.23	0.73
Implicit Prompts	-0.03	0.98	0.04	0.98	0.04	0.97
SST-2	0.09	1	0.06	0.99	0.05	0.98
SemEval-2019	0.38	0.8	0.26	0.91	0.17	0.89

Table 2: Kappa agreement between relevance of Human-A and Gpt-4o-mini ($\mathcal{K}_{H,M}^{rel}$) and corresponding Inclusion score (\mathcal{I}_{HCM}) for increasing k .

the LLMs in labeling a demonstration as relevant. To compensate for this, we use two types of relevance scores (Human and LLM) in the following sections for our correlation analyses. For the LLM annotator, which scales better, we also generate the relevance scores for *all* test data prompts for $k = 1, 10,$ and 20 examples.

3.3.2 Correlation between Example Relevance and Model Similarity

Based on prior discussions, we know that ICL is similar to classifiers in many situations; the next step is to identify the conditions under which this similarity occurs. To explore this, we examine how the relevance scores of humans and LLMs are correlated with the Kappa agreement between models. For simplicity, we refer to the Kappa agreement between the labels predicted by the classifiers and LLM as model Kappa, denoted by $\mathcal{K}_{C,L}^{model}$.

To measure the correlation between model similarity and example relevance, we employ two types of correlations that are widely used to measure the association between two variables: the Pearson correlation coefficient (r) and the R-squared (R^2) coefficient (Sedgwick, 2012; Wright, 1921). r captures the linear correlation between two variables by computing the ratio between the covariance of two variables and the product of their standard deviations; it ranges between -1 and 1 . R^2 is the proportion of the variation in one variable that is predictable from the other one. For example, we interpret the R^2 between example similarity and the Kappa agreement between two models as the fraction of the Kappa score that can be explained by example relevance.

Formally defining, we measure the correlation between the $\mathcal{K}_{C,L}^{model}$ and two variables: the first one is the average relevance score determined by a human (denoted by \mathcal{R}_{Human}), and the second is by the respective LLM as an annotator (denoted

Datasets \ Neighbors	$k = 1$			$k = 10$			$k = 20$		
	\mathcal{R}_{Human}	$\mathcal{K}_{C,L}^{model}$	\mathcal{R}_{LLM}	\mathcal{R}_{Human}	$\mathcal{K}_{C,L}^{model}$	\mathcal{R}_{LLM}	\mathcal{R}_{Human}	$\mathcal{K}_{C,L}^{model}$	\mathcal{R}_{LLM}
AG News	0.9	0.87	1	0.794	0.87	0.976	0.7	0.87	0.952
News Category	0.8	0.87	0.88	0.664	0.87	0.834	0.553	0.84	0.795
AuTexTification	0.5	0.06	0.66	0.312	0.03	0.546	0.267	0.02	0.507
Implicit Prompts	0.96	0.19	0.98	0.756	0.39	0.966	0.627	0.54	0.956
SST-2	0.7	0.31	0.98	0.596	0.49	0.97	0.471	0.57	0.961
SemEval-2019	0.2	0.12	0.38	0.21	0.3	0.57	0.154	0.34	0.581
<i>Correlation-A</i>	$r: 0.559$ $R^2: 0.313$			$r: 0.728$ $R^2: 0.530$			$r: 0.795$ $R^2: 0.631$		
<i>Correlation-B</i>	$r: 0.528$ $R^2: 0.279$			$r: 0.668$ $R^2: 0.446$			$r: 0.768$ $R^2: 0.589$		

Table 3: Correlation comparison with **GPT-4o-mini** and **kNN** on 50 data prompts. The top part of the table shows three columns: the human-annotated relevance score for ICL examples (\mathcal{R}_{Human}); Kappa agreement between the labels of classifier and LLM ($\mathcal{K}_{C,L}^{model}$); and the machine-annotated relevance score on ICL examples used (\mathcal{R}_{LLM}). The bottom part shows two correlations; *Correlation-A* = $Corr(\mathcal{K}_{C,L}^{model}, \mathcal{R}_{Human})$ (in green) and *Correlation-B* = $Corr(\mathcal{K}_{C,L}^{model}, \mathcal{R}_{LLM})$ (in purple).

Datasets \ Neighbors	$k = 1$		$k = 10$		$k = 20$	
	$\mathcal{K}_{C,L}^{model}$	\mathcal{R}_{LLM}	$\mathcal{K}_{C,L}^{model}$	\mathcal{R}_{LLM}	$\mathcal{K}_{C,L}^{model}$	\mathcal{R}_{LLM}
AG News	0.82	0.974	0.87	0.934	0.86	0.917
News Category	0.67	0.924	0.67	0.839	0.69	0.792
AuTexTification	0.16	0.608	0.09	0.511	-0.02	0.468
Implicit Prompts	0.3	0.868	0.33	0.866	0.28	0.861
SST-2	0.37	0.944	0.44	0.929	0.46	0.911
SemEval-2019	0.17	0.863	0.3	0.868	0.32	0.87
<i>Correlation-B</i>	$r: 0.674$ $R^2: 0.455$		$r: 0.676$ $R^2: 0.457$		$r: 0.691$ $R^2: 0.477$	

Table 4: Correlation comparison with **GPT-4o-mini** and **kNN** on *all* data prompts. The table shows the correlation between: Kappa agreement between the outputs of classifier(kNN) and LLM ($\mathcal{K}_{C,L}^{model}$); and the annotated relevance score by LLM on ICL examples (\mathcal{R}_{LLM}).

by \mathcal{R}_{LLM}). Intuitively, we hypothesize that there should be a direct correlation between both relevance scores and model Kappa. In other words, higher example relevance should correspond to a higher Kappa agreement score.¹⁰

Building on this hypothesis, we observe three key findings in Table 3. First, the Pearson Correlation (r) values are positive for each correlation, confirming that the more relevant the examples are, the more similar the LLM and classifier behave. This holds for both types of correlations, *A* and *B*, as defined in the table. Second, we find that increasing the number of examples (k) per test data point, the data more strongly supports these observations, suggesting that the LLM increasingly relies on ICL examples as more instances are provided in the context. Third, focusing on the R^2

¹⁰Since most of the datasets involve multi-class classification, we use the multi-class Kappa agreement formula (Cohen, 1960).

value for *Correlation-A* for $k = 20$, we indicate that up to 63.1% of the behavior similarity can be explained by example relevance.¹¹

3.3.3 Scaling Up the Correlation Analysis to the Entire Test Partitions

As mentioned earlier, manual relevance is limited and expensive; hence, to scale up the experiment, we calculate the average relevance score on the annotations given by the underlying LLMs. This implies that the same procedure is extended to all test prompts, with the three LLMs employed as annotators for the relevance score of examples.

Consistent with the observations in Table 3, Table 4 shows that LLM maintains high relevance scores (\mathcal{R}_{LLM}) across k values because of high-quality semantically similar examples. Further-

¹¹Tables 10 through 14 in Appendix I show the results of all the analyses of correlation for the three LLMs investigated with classifiers.

more, the correlation value exhibits an upward trend, reflecting a similar pattern to that previously observed with both human and machine relevance on 50 test prompts. Note that in some cases, there is a slight drop in relevance scores from $k = 10$ to $k = 20$, which occurs because an increase in the number of examples introduces low-quality examples in the set, thereby reducing the overall relevance score. Because this drop aligns with model Kappa in these cases, the correlation values are not affected and exhibit the same trend as before.

4 Discussion

We hypothesize that when the relevance of the ICL examples is high, both LLMs and the classifiers rely on them for their predictions. When example relevance is low, differences begin to emerge; LLMs can back off to their parametric memory and still perform well in some cases. In contrast, the classifiers’ performance is directly impacted by example relevance. Additionally, LLM indeed backs off to its parametric memory when using low-relevance demonstrations in ICL.¹²

We summarize the key observations from this study in Algorithm 1. The algorithm indicates that when the examples selected to be used by the classifier or ICL in LLMs are relevant (which was established using both human and LLM decisions), ICL performs and behaves similarly to kNN (and a little less so to LR). We validate this observation by showing that there is a high correlation between example relevance and similarity between ICL and the classifier. We also notice that when the relevance of examples is low, the two models diverge: the performance of classifiers generally drops, whereas the performance of the LLM is contingent on its pre-training and parametric memory, i.e., if the task at hand is well represented in its parametric memory, it continues to perform well.

Similar to our analysis, the task recognition method of Pan et al. (2023) also makes the distinction between non-parametric and parametric memory. However, we explain why both are important, and how ICL may choose one or another depending on the relevance of the demonstrations. “Task recognition” encourages LLMs to back off to a certain part of the parametric memory, so we think that their observation complements nicely our own.

¹²This was validated by comparing the outputs of the LLM with poor-quality examples with a zero-shot configuration of the same LLM. Please see Appendix J.

Algorithm 1 Summary of LLM with ICL behavior

```

 $top_k \leftarrow$  top  $k$  semantically similar examples
if  $top_k$  is relevant then
    LLMs  $\approx$  kNN
else if  $top_k$  is less relevant then
    Performance of all classifiers drops
    Performance of LLMs may remain high, possibly
    due to parametric memory information
end if

```

This summary suggests an important practical insight: in the situations where the relevance of the ICL examples is high, the LLM can be safely replaced by the simpler kNN algorithm, which can be efficiently implemented using vector databases such as FAISS (Douze et al., 2024). This, of course, puts the onus on the methods for the example selection, but this is research that has already shown progress (Akyürek et al., 2023; Zhou et al., 2025; Shen et al., 2024; Fu et al., 2024; Nafar et al., 2025; Wang et al., 2024; An et al., 2023).

5 Conclusion

In this paper, we examine how closely ICL behaves to classifiers based on kNN or GD trained on the same ICL examples. We find that ICL behaves more similarly to kNN when subjected to certain conditions driven by the relevance of examples. This suggests that if one can strategically select highly relevant examples, they can just use kNN. However, in cases where the examples are of lower relevance, LLMs demonstrate greater robustness to noise due to their ability to back off to their parametric memory. As a result, LLMs become a more suitable choice in such scenarios, as prior knowledge allows them to perform better.

6 Acknowledgments

Mihai Surdeanu declares a financial interest in lum.ai. This interest has been properly disclosed to the University of Arizona Institutional Review Committee and is managed in accordance with its conflict of interest policies.

7 Limitations

In our current work, we focus on three LLMs and six text classification datasets, which provide initial insights into the observed patterns. We believe our observations are supported empirically by our current work and are already actionable. However, to strengthen our understanding of the underlying behavior of ICL, we plan to incorporate more LLMs,

multiple datasets, and more NLP tasks in future work. This will help validate our conclusions with greater confidence, ensuring that our results hold across diverse scenarios.

Further, our comparison between kNN and gradient descent and ICL used a linear GD-based classifier (logistic regression). This is a common strategy in probing (Hewitt and Liang, 2019; Zhao and Bethard, 2020). However, it is possible that ICL behaves closer to a non-linear GD-based classifier, such as a feed-forward neural network.

One concern is whether there is a causal link between the two variables that are correlated. i.e., the relevance scores of demonstrations and the Kappa agreement between LLMs and classifiers. In future work, we propose to compare two identical configurations (i.e., same LLM, same dataset, same number of demonstrations k), with the only difference being that one has access to high-quality demonstrations while the other uses low-quality examples. If causality holds, the Kappa values would be high for all configurations, and they should drop when we use bad examples.

8 Ethical Considerations

This work involves some datasets that may contain offensive content, which could be distressing to some readers. We have taken measures by using appropriate disclaimers to acknowledge the offensive language.

References

- Kwangjun Ahn, Xiang Cheng, Hadi Daneshmand, and Suvrit Sra. 2023. [Transformers learn to implement preconditioned gradient descent for in-context learning](#). *Preprint*, arXiv:2306.00297.
- AI@Meta. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Ekin Akyürek, Dale Schuurmans, Jacob Andreas, Tengyu Ma, and Denny Zhou. 2023. [What learning algorithm is in-context learning? investigations with linear models](#). *Preprint*, arXiv:2211.15661.
- Afra Amini and Massimiliano Ciaramita. 2023. [In-context probing: Toward building robust classifiers via probing large language models](#). *Preprint*, arXiv:2305.14171.
- Shengnan An, Bo Zhou, Zeqi Lin, Qiang Fu, Bei Chen, Nanning Zheng, Weizhu Chen, and Jian-Guang Lou. 2023. [Skill-based few-shot selection for in-context learning](#). *Preprint*, arXiv:2305.14210.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D 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 Ziegler, Jeffrey Wu, Clemens Winter, Chris 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*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1):37–46.
- Damai Dai, Yutao Sun, Li Dong, Yaru Hao, Shuming Ma, Zhifang Sui, and Furu Wei. 2023. [Why can gpt learn in-context? language models implicitly perform gradient descent as meta-optimizers](#). *Preprint*, arXiv:2212.10559.
- Amit Das, Mostafa Rahgouy, Dongji Feng, Zheng Zhang, Tathagata Bhattacharya, Nilanjana Raychawdhary, Fatemeh Jamshidi, Vinija Jain, Aman Chadha, Mary J. Sandage, Lauramarie Pope, Gerry V. Dozier, and Cheryl D. Seals. 2024. [Offensivelang: A community-based implicit offensive language dataset](#). *IEEE Access*, 12:185661–185672.
- Gilad Deutch, Nadav Magar, Tomer Bar Natan, and Guy Dar. 2024. [In-context learning and gradient descent revisited](#). *Preprint*, arXiv:2311.07772.
- Matthijs Douze, Alexandr Guzhva, Chengqi Deng, Jeff Johnson, Gergely Szilvassy, Pierre-Emmanuel Mazaré, Maria Lomeli, Lucas Hosseini, and Hervé Jégou. 2024. [The faiss library](#).
- Sahibsingh A. Dudani. 1976. [The distance-weighted k-nearest-neighbor rule](#). *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-6(4):325–327.
- E. Fix and J.L. Hodges. 1951. [Discriminatory Analysis: Nonparametric Discrimination: Consistency Properties](#). USAF School of Aviation Medicine.
- Deqing Fu, Tian-Qi Chen, Robin Jia, and Vatsal Sharan. 2024. [Transformers learn to achieve second-order convergence rates for in-context linear regression](#). *Preprint*, arXiv:2310.17086.
- John Hewitt and Percy Liang. 2019. [Designing and interpreting probes with control tasks](#). *Preprint*, arXiv:1909.03368.
- Michael P LaValley. 2008. Logistic regression. *Circulation*, 117(18):2395–2399.
- Katerina Margatina, Timo Schick, Nikolaos Aletras, and Jane Dwivedi-Yu. 2023. [Active learning principles for in-context learning with large language models](#). *Preprint*, arXiv:2305.14264.

- 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?](#) *Preprint*, arXiv:2202.12837.
- Rishabh Misra. 2022. News category dataset. *arXiv preprint arXiv:2209.11429*.
- Aliakbar Nafar, K. Brent Venable, and Parisa Kordjamshidi. 2025. [Learning vs retrieval: The role of in-context examples in regression with large language models](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 8206–8229, Albuquerque, New Mexico. Association for Computational Linguistics.
- OpenAI. 2024. [Gpt-4o system card](#). *Preprint*, arXiv:2410.21276.
- Jane Pan, Tianyu Gao, Howard Chen, and Danqi Chen. 2023. [What in-context learning "learns" in-context: Disentangling task recognition and task learning](#). *Preprint*, arXiv:2305.09731.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-bert: Sentence embeddings using siamese bert-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Areg Mikael Sarvazyan, José Ángel González, Marc Franco-Salvador, Francisco Rangel, Berta Chulvi, and Paolo Rosso. 2023. [Overview of autextification at iberlef 2023: Detection and attribution of machine-generated text in multiple domains](#). *Preprint*, arXiv:2309.11285.
- Philip Sedgwick. 2012. Pearson’s correlation coefficient. *Bmj*, 345.
- Lingfeng Shen, Aayush Mishra, and Daniel Khashabi. 2024. [Do pretrained transformers learn in-context by gradient descent?](#) *Preprint*, arXiv:2310.08540.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Johannes Von Oswald, Eyvind Niklasson, Ettore Randazzo, Joao Sacramento, Alexander Mordvintsev, Andrey Zhmoginov, and Max Vladymyrov. 2023. [Transformers learn in-context by gradient descent](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 35151–35174. PMLR.
- Johannes von Oswald, Maximilian Schlegel, Alexander Meulemans, Seijin Kobayashi, Eyvind Niklasson, Nicolas Zucchet, Nino Scherrer, Nolan Miller, Mark Sandler, Blaise Agüera y Arcas, Max Vladymyrov, Razvan Pascanu, and João Sacramento. 2024. [Uncovering mesa-optimization algorithms in transformers](#). *Preprint*, arXiv:2309.05858.
- Xinyi Wang, Wanrong Zhu, Michael Saxon, Mark Steyvers, and William Yang Wang. 2024. [Large language models are latent variable models: Explaining and finding good demonstrations for in-context learning](#). *Preprint*, arXiv:2301.11916.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022. [Emergent abilities of large language models](#). *Preprint*, arXiv:2206.07682.
- Sewall Wright. 1921. Correlation and causation. *Journal of agricultural research*, 20(7):557.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jianxin Yang, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Xuejing Liu, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, Zhifang Guo, and Zhihao Fan. 2024. [Qwen2 technical report](#). *Preprint*, arXiv:2407.10671.
- Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019. [SemEval-2019 task 6: Identifying and categorizing offensive language in social media \(OffensEval\)](#). In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 75–86, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. [Character-level convolutional networks for text classification](#). In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.
- Yiming Zhang, Shi Feng, and Chenhao Tan. 2022. [Active example selection for in-context learning](#). *Preprint*, arXiv:2211.04486.
- Yiyun Zhao and Steven Bethard. 2020. [How does BERT’s attention change when you fine-tune? an analysis methodology and a case study in negation scope](#). In *Proceedings of the 58th Annual Meeting of*

the Association for Computational Linguistics, pages 4729–4747, Online. Association for Computational Linguistics.

Zhenghao Zhou, Robert Frank, and R. Thomas McCoy. 2025. [Is in-context learning a type of error-driven learning? evidence from the inverse frequency effect in structural priming](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 11712–11725, Albuquerque, New Mexico. Association for Computational Linguistics.

A Analysis of Manual Annotation Quality

Table 5 presents the inter-annotator agreement for the manual annotations of example relevance. The Kappa value, representing the level of agreement between annotators, was computed based on the manual annotations of two data points per dataset for $k = 10$ and $k = 20$. In total, $2 \times (10 + 20) = 60$ individual examples per dataset were annotated by each annotator.

Nearest Neighbors	Kappa
$k = 10$	0.70
$k = 20$	0.68

Table 5: Kappa agreement scores between the two annotators across all datasets.

These values indicate substantial agreement between the two human annotators, which provides evidence that the manual annotations of example relevance are reliable.

B Dataset Labels & Statistics

The complete statistics of all datasets used in the experiments are presented in Table 6.

Datasets	Train	Total Test	Sampled Test	#Labels
AG News	120000	7600	1001	4
News Category	40000	10000	1000	10
AuTexTification	33845	21832	1000	2
Implicit Prompts	6615	1655	1655	2
SST-2	67349	872	872	2
SemEval-2019	74682	1000	1000	4

Table 6: Data statistics of all the datasets used.

Additionally, the labels for the six datasets used in this work are presented in Table 7. We employ both binary and multi-class labels to ensure the generality of our approach.

Datasets	Labels
AG News	World, Sports, Business, Sci/Tech
News Category	Wellness, Politics, Entertainment, Travel, Style & Beauty, Parenting, Food & Drink, World News, Business, Sports
AuTexTification	Human, Generated
Implicit Prompts	Offensive, Not Offensive
SST-2	Positive, Negative
SemEval-2019	Negative, Positive, Neutral, Irrelevant

Table 7: Labels for the six datasets used in this work.

C Logistic Regression Model Training Using Top-k Nearest Examples

The classifier in this work is a vanilla logistic regression, similar to probing classifiers, where we have a linear layer on top of contextualized representations (Amini and Ciaramita, 2023).

We train a logistic regression model on ICL examples using their text and labels. The examples are first converted into numerical representations using a TfidfVectorizer vectorizer, which transforms each sentence into a high-dimensional feature vector. An LR model is trained only on these top- k examples as features with their labels. The same vectorizer is used to transform the test prompt into the same feature space.

D Example Prompt Used for Text Classification in ICL

The following prompt is used in the case of GPT-4o-mini for prompt engineering corresponding to AG News.

```
{"role": "system",
"content": f"Example: We know that
the classification for the text
'{examples[i]}', we have
answer: '{labels[i]}'."}

```

```
{"role": "system",
"content": f"According to the
above provided examples, classify
the following text. Answer as
World, Sports, Business, Sci/Tech
with no explanation."}

```

```
{"role": "user", "content": text}

```


Dataset	Test Data Point	Nearest Neighbors	Relevance Score(0/1)
AG News	Airlines Agree to Cuts at O’Hare	– Airlines Agree to Cuts at O’Hare	1
		– Airlines to Cut Flights at Chicago O’Hare	1
		– US airlines agree to cut flights at Chicago #39;s O #39;Hare	1
		– 2 Big Carriers at O #39;Hare to Cut Flights	1
		– 2 Big Carriers at O #39;Hare to Cut Flights	1
		– UPDATE 3-US airlines agree to cut flights at Chicago #39;s O #39;Hare	1
		– Airlines agree to limit O #39;Hare arrivals	1
		– FAA: Flight-Reduction Deal Set for O #39;Hare	1
		– O #39;Hare to reduce flight arrivals	1
		– More Job Cuts Likely At American Airlines	0
SemEval-2019 (ApexLegends)	Bout to fuck around and stream. @PlayApex grind and @Brawlhalla afterwards. mixer.com/ShinobiSZN #BitGang https://t.co/aBUpN6yjfV	– My goal is just a funny man lmaoo	0
		– said My aim is just ridiculous man lmaoo	0
		– My aim is just be man lmaoo	0
		– My Wife is just ridiculous man lmaoo	0
		– @PlayApex Past 10 days of pub matchmaking has been a nightmare. Not sure what changed but noticeably worse team mates.	0
		– @PlayApex servers are expanding. 200d.	0
		– @PlayApex servers are whack. The [UNK]..	0
		– . . Stupid robot. pic.twitter.com/nX5ZFPM85G	0
		– @PlayApex so you guys have zero humans on your life? Best friends account was CLEARLY hacked for illegal purpose of aimbotting/hacking in game, he sent 2 tickets and never got a human response. He has proof, those guys won’t even look at it. We have supported your business since day	0
		– @ PlayApex so you guys don’t have people in customer support? The best friends account was CLEARLY hacked to hack / hack into the game, it sent 2 tickets and never got a human response.	0

Table 8: **WARNING: This table contains language that might be offensive to some.** 10 examples corresponding to a given test data point from two datasets. The first row indicates a test data point associated with high-quality examples (relevance score: 9/10), whereas the second row presents a test data point with low-quality examples (relevance score: 0/10).

In the given prompt, we first provide a set of example–label pairs determined by the value of k based on their similarity scores (top_k in Algorithm 1). In the next part of the prompt, the model is asked to generate the label for the provided test text.

E Relevance Score Labeling of Examples

We mark the top- k high-quality examples with a Boolean label for each example to indicate its relevance. We assign a value of ‘1’ to the examples that are relevant to the input test data prompt, and ‘0’ to those that are not. Table 8 shows two cases where one dataset contains more relevant examples than the other.

F Additional Contingency Matrices for Unweighted kNN

As discussed in the paper, we have generated contingency matrices to validate the classification of labels for unweighted-kNN and LR with respect to three LLMs (refer to Section 3.2.2). Figures 5 to 8 show the contingency matrices for all other

datasets. As previously noted, for several datasets, we observe that the classifiers closely align with the classifications produced by LLMs, with only minor differences as k increases.

By analyzing the figures, we observe that, except for AG News (Figure 3) and News Category (Figures 4), the other datasets exhibit lower diagonal dominance. This is likely because the performance of the classifiers heavily depends on the examples, and on such datasets, the classifiers may be influenced by the way the examples are labeled.

This is evident because, although the performance metrics of the classifiers are comparable to those of LLMs in the case of AuTexTification (as shown in Table 1), the contingency matrices reveal inconsistencies in label classification (Figure 5). Similarly, in the Implicit Prompts (Figure 6), the dataset is designed to contain implicit offensive prompts, which means most examples are labeled as ‘Offensive’. Consequently, the classifiers exhibit strong alignment with these labels.

The SST-2 (Figure 7) and SemEval-2019 (Figure 8) datasets show moderate diagonal dominance, meaning that correct predictions are somewhat

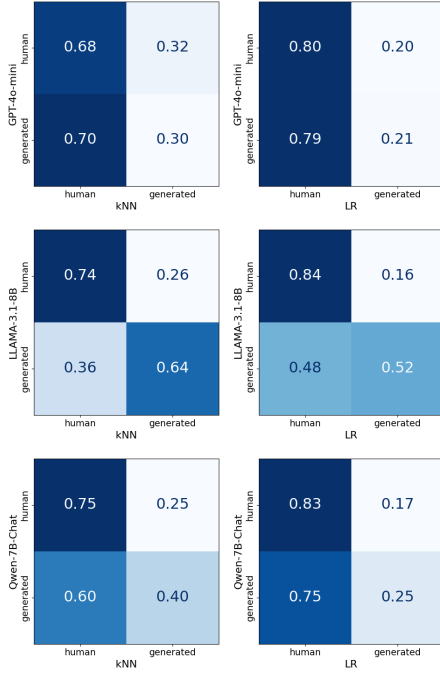


Figure 5: Contingency matrices for AuTexTification: The same configuration follows from Figure 3.

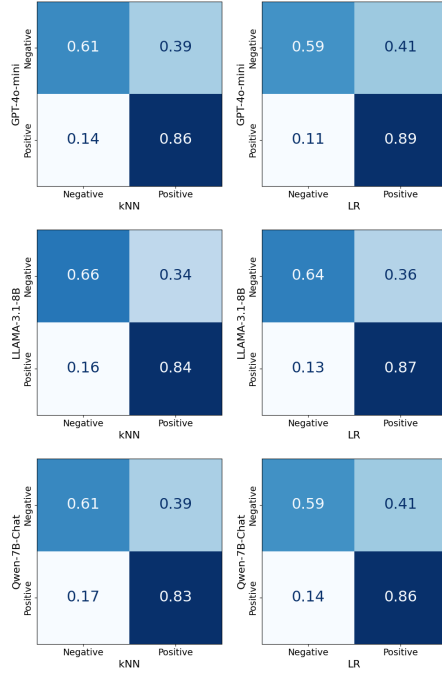


Figure 7: Contingency matrices for SST-2: The same configuration follows from Figure 3.

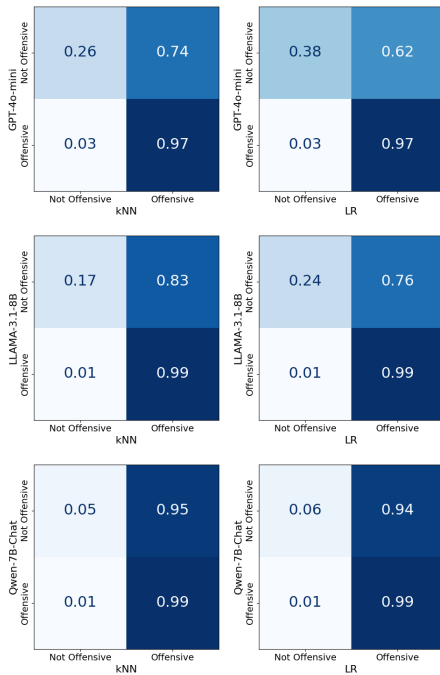


Figure 6: Contingency matrices for Implicit Prompts: The same configuration follows from Figure 3.

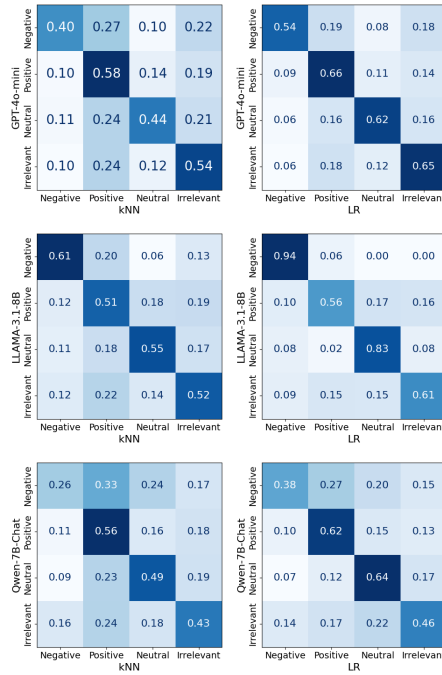


Figure 8: Contingency matrices for SemEval-2019: The same configuration follows from Figure 3.

G Weighted kNN Performance Evaluation

more frequent along the main diagonal of the confusion matrix. For these datasets, classifiers generally perform worse than LLMs, as indicated in Table 1.

Overall, these analyses suggest that classifiers and LLMs overlap in both performance and behavior in some cases, but not all.

A reasonable question that arises from the results in Section 3.2 is: *Why not compare model performance when the models are informed about the example similarity scores?* To explore this, we aim to examine how similar the behavior of ICL

Dataset	$k = 1$				$k = 10$				$k = 20$				$k = 30$			
	kNN	GPT	Llama	Qwen	kNN	GPT	Llama	Qwen	kNN	GPT	Llama	Qwen	kNN	GPT	Llama	Qwen
AG News	0.85	0.86	0.82	0.81	0.88	0.89	0.88	0.88	0.88	0.89	0.88	0.88	0.87	0.90	0.87	0.87
News Category	0.80	0.86	0.80	0.77	0.83	0.88	0.85	0.81	0.83	0.88	0.85	0.81	0.83	0.88	0.85	0.80
AuTexTification	0.54	0.55	0.54	0.53	0.59	0.57	0.60	0.54	0.59	0.57	0.60	0.54	0.58	0.57	0.58	0.54
Implicit Prompts	0.69	0.72	0.69	0.70	0.78	0.77	0.75	0.71	0.78	0.77	0.75	0.71	0.79	0.77	0.75	0.71
SST-2	0.68	0.93	0.87	0.91	0.73	0.94	0.91	0.93	0.73	0.94	0.91	0.93	0.72	0.95	0.92	0.93
SemEval-2019	0.39	0.58	0.54	0.46	0.49	0.73	0.65	0.56	0.49	0.73	0.65	0.56	0.49	0.73	0.64	0.55

Table 9: Test accuracy on six datasets for weighted-kNN and three LLMs with weighted ICL, using 1, 10, 20, and 30 examples.

would be to weighted kNN if the ICL prompt also includes information about the similarity scores of examples.¹³ This analysis is built on the idea of conducting a comparative evaluation between weighted ICL and weighted kNN.¹⁴

We report the performance values of weighted kNN with three LLMs, considering all datasets for different numbers of k examples in Table 9. As noted earlier in Section 3.1.1, the performance is nearly identical to that of the unweighted kNN, suggesting that the inclusion of similarity scores does not substantially influence the results, and the

¹³For simplicity, we call the scenario as *weighted ICL*, where we add the similarity scores of each example to the prompt of the original ICL, given in Appendix D.

¹⁴Definition of weighted kNN is discussed in 3.1.1.




Figure 9: Contingency matrices for AG News: Comparison of weighted kNN vs GPT (top), weighted kNN vs Llama (middle), and weighted kNN vs Qwen (bottom).

behavior of ICL and kNN remains the same.

In line with the observations from Table 1, Table 9 indicates that the accuracy of weighted kNN and the weighted ICL on three LLMs is comparable across most cases. Additionally, the overall trend of accuracy values remains consistent with Table 1.

To further verify whether the performance of the models is indeed aligned, we examine the contingency matrices between the predictions of each LLM with weighted kNN. Figures 9 to 14 present the contingency matrices between weighted kNN (X-axis) and the three LLMs (Y-axis) on the six datasets.




Figure 10: Contingency matrices for News Category: The same configuration follows from Figure 9.




Figure 11: Contingency matrices for AuTexTification: The same configuration follows from Figure 9.




Figure 13: Contingency matrices for SST-2: The same configuration follows from Figure 9.




Figure 12: Contingency matrices for Implicit Prompts: The same configuration follows from Figure 9.




Figure 14: Contingency matrices for SemEval-2019: The same configuration follows from Figure 9.

The resulting patterns of contingency matrices for weighted closely resemble those obtained for the unweighted variant of kNN, indicating consistent trends across both kNN configurations.

H Does ICL Behave Closer to kNN or LR?

The Kappa scores reported in the six tables (3, 10, and 11 through 14), which measure behavior similarity between kNN/LR and the LLMs, show small but consistently higher values for kNN. For exam-

ple, for GPT-4o-mini with $k = 10$, the kNN model achieves a model Kappa ($\mathcal{K}_{C,L}^{model}$) value of 0.87 on AG News (first row in Table 3), whereas Logistic Regression (LR) attains a Kappa of 0.84 (first row in Table 10). On average, this observation holds for the other datasets and LLMs.

Further, the correlation values in these tables show a more intuitive behavior for kNN. This is evident from the increasing Pearson (r) and R^2 values as k grows, e.g., R^2 grows from 31.3% to 63.1% from $k = 1$ to $k = 20$ for *Correlation-A* (Table 3). This trend does not hold for LR: R^2 drops from 31.3% to 26.8% in the same scenario (Table 10).

These observations suggest that ICL behaves closer to kNN than LR. However, as we have a very simple gradient-based classifier, the results may vary in larger models. In our setting involving simplified models, we focus on understanding the behavior of ICL with these classifiers.

I Additional Correlation Results

Extending our analysis from Table 3 to the comparison between GPT-4o-mini and LR, we present the correlation values in Table 10. Evidently, we note that R^2 value for *Correlation-A* with $k = 20$ indicates that only up to 26.8% of the behavioral similarity can be explained by example relevance. This demonstrates that KNN exhibits a stronger correlation compared to LR.

Additionally, Tables 11 and 12 show the correlation results for Llama-3.1-8B with kNN and LR, respectively. As expected, kNN has a comparatively better correlation with the LLM when compared to LR. Similarly, we extend our analysis for QWEN-7B-Chat in Tables 13 and 14. kNN tends to perform closer to the LLM in this case as well.

In summary, the observations from Section 3.3.2 generally hold for the other two LLMs, with a couple of caveats. First, the correlations are weaker for $k = 1$, suggesting that one ICL example is not sufficient to influence the predictions of the two

Datasets \ Neighbors	$k = 1$			$k = 10$			$k = 20$		
	\mathcal{R}_{Human}	$\mathcal{K}_{C,L}^{model}$	\mathcal{R}_{LLM}	\mathcal{R}_{Human}	$\mathcal{K}_{C,L}^{model}$	\mathcal{R}_{LLM}	\mathcal{R}_{Human}	$\mathcal{K}_{C,L}^{model}$	\mathcal{R}_{LLM}
AG News	0.9	0.87	1	0.794	0.84	0.976	0.7	0.84	0.952
News Category	0.8	0.87	0.88	0.664	0.84	0.834	0.553	0.87	0.795
AuTextification	0.5	0.06	0.66	0.312	0.11	0.546	0.267	0.05	0.507
Implicit Prompts	0.96	0.19	0.98	0.756	0.19	0.966	0.627	0.3	0.956
SST-2	0.7	0.31	0.98	0.596	0.42	0.97	0.471	0.53	0.961
SemEval-2019	0.2	0.12	0.38	0.21	0.46	0.57	0.154	0.49	0.581
<i>Correlation-A</i>	$r: 0.559$ $R^2: 0.313$			$r: 0.431$ $R^2: 0.186$			$r: 0.517$ $R^2: 0.268$		
<i>Correlation-B</i>	$r: 0.528$ $R^2: 0.279$			$r: 0.368$ $R^2: 0.135$			$r: 0.494$ $R^2: 0.244$		

Table 10: Comparison of **GPT-4o-mini** and **LR**. The same configuration follows from Table 3.

Datasets \ Neighbors	$k = 1$			$k = 10$			$k = 20$		
	\mathcal{R}_{Human}	$\mathcal{K}_{C,L}^{model}$	\mathcal{R}_{LLM}	\mathcal{R}_{Human}	$\mathcal{K}_{C,L}^{model}$	\mathcal{R}_{LLM}	\mathcal{R}_{Human}	$\mathcal{K}_{C,L}^{model}$	\mathcal{R}_{LLM}
AG News	0.9	0.92	0.86	0.794	1	0.722	0.7	0.89	0.696
News Category	0.8	0.89	0.82	0.664	0.89	0.762	0.553	0.91	0.736
AuTextification	0.5	1	0.8	0.312	0.67	0.764	0.267	0.35	0.749
Implicit Prompts	0.96	0.29	0.8	0.756	0.52	0.672	0.627	0.21	0.625
SST-2	0.7	0.51	0.92	0.596	0.74	0.884	0.471	0.7	0.868
SemEval-2019	0.2	0.15	0.48	0.21	0.31	0.594	0.154	0.32	0.594
<i>Correlation-A</i>	$r: 0.315$ $R^2: 0.099$			$r: 0.665$ $R^2: 0.442$			$r: 0.516$ $R^2: 0.266$		
<i>Correlation-B</i>	$r: 0.582$ $R^2: 0.339$			$r: 0.594$ $R^2: 0.352$			$r: 0.503$ $R^2: 0.253$		

Table 11: Comparison of **Llama-3.1-8B** and **kNN**. The same configuration follows from Table 3.

Datasets \ Neighbors	$k = 1$			$k = 10$			$k = 20$		
	\mathcal{R}_{Human}	$\mathcal{K}_{C,L}^{model}$	\mathcal{R}_{LLM}	\mathcal{R}_{Human}	$\mathcal{K}_{C,L}^{model}$	\mathcal{R}_{LLM}	\mathcal{R}_{Human}	$\mathcal{K}_{C,L}^{model}$	\mathcal{R}_{LLM}
AG News	0.9	0.92	0.86	0.794	0.92	0.722	0.7	0.87	0.696
News Category	0.8	0.89	0.82	0.664	0.89	0.762	0.553	0.93	0.736
AuTexTification	0.5	1	0.8	0.312	0.62	0.764	0.267	0.46	0.749
Implicit Prompts	0.96	0.29	0.8	0.756	0.41	0.672	0.627	0.26	0.625
SST-2	0.7	0.51	0.92	0.596	0.65	0.884	0.471	0.74	0.868
SemEval-2019	0.2	0.15	0.48	0.21	0.43	0.594	0.154	0.39	0.594
<i>Correlation-A</i>	$r: 0.315$ $R^2: 0.099$			$r: 0.480$ $R^2: 0.230$			$r: 0.449$ $R^2: 0.202$		
<i>Correlation-B</i>	$r: 0.582$ $R^2: 0.339$			$r: 0.453$ $R^2: 0.206$			$r: 0.552$ $R^2: 0.305$		

Table 12: Comparison of **Llama-3.1-8B** and **LR**. The same configuration follows from Table 3.

Datasets \ Neighbors	$k = 1$			$k = 10$			$k = 20$		
	\mathcal{R}_{Human}	$\mathcal{K}_{C,L}^{model}$	\mathcal{R}_{LLM}	\mathcal{R}_{Human}	$\mathcal{K}_{C,L}^{model}$	\mathcal{R}_{LLM}	\mathcal{R}_{Human}	$\mathcal{K}_{C,L}^{model}$	\mathcal{R}_{LLM}
AG News	0.9	0.92	0.82	0.794	0.89	0.688	0.7	0.89	0.629
News Category	0.8	0.37	0.54	0.664	0.63	0.418	0.553	0.76	0.389
AuTexTification	0.5	0.02	0.6	0.312	-0.12	0.586	0.267	0.1	0.549
Implicit Prompts	0.96	0.11	0.8	0.756	0.15	0.694	0.627	0.1	0.646
SST-2	0.7	0.34	0.68	0.596	0.49	0.766	0.471	0.61	0.762
SemEval-2019	0.2	0.08	0.42	0.21	0.15	0.474	0.154	0.33	0.476
<i>Correlation-A</i>	$r: 0.525$ $R^2: 0.276$			$r: 0.691$ $R^2: 0.477$			$r: 0.506$ $R^2: 0.256$		
<i>Correlation-B</i>	$r: 0.526$ $R^2: 0.277$			$r: 0.128$ $R^2: 0.016$			$r: -0.035$ $R^2: 0.001$		

Table 13: Comparison of **QWEN-7B-Chat** and **kNN**. The same configuration follows from Table 3.

Datasets \ Neighbors	$k = 1$			$k = 10$			$k = 20$		
	\mathcal{R}_{Human}	$\mathcal{K}_{C,L}^{model}$	\mathcal{R}_{LLM}	\mathcal{R}_{Human}	$\mathcal{K}_{C,L}^{model}$	\mathcal{R}_{LLM}	\mathcal{R}_{Human}	$\mathcal{K}_{C,L}^{model}$	\mathcal{R}_{LLM}
AG News	0.9	0.92	0.82	0.794	0.87	0.688	0.7	0.87	0.629
News Category	0.8	0.37	0.54	0.664	0.61	0.418	0.553	0.76	0.389
AuTexTification	0.5	0.02	0.6	0.312	-0.14	0.586	0.267	0.02	0.549
Implicit Prompts	0.96	0.11	0.8	0.756	0.11	0.694	0.627	0.06	0.646
SST-2	0.7	0.34	0.68	0.596	0.42	0.766	0.471	0.57	0.762
SemEval-2019	0.2	0.08	0.42	0.21	0.21	0.474	0.154	0.38	0.476
<i>Correlation-A</i>	$r: 0.525$ $R^2: 0.276$			$r: 0.629$ $R^2: 0.395$			$r: 0.453$ $R^2: 0.205$		
<i>Correlation-B</i>	$r: 0.526$ $R^2: 0.277$			$r: 0.049$ $R^2: 0.002$			$r: -0.094$ $R^2: 0.009$		

Table 14: Comparison of **QWEN-7B-Chat** and **LR**. The same configuration follows from Table 3.

LLMs. Second, the correlation values increase considerably from $k = 1$ to $k = 10$, but not always from $k = 10$ to $k = 20$. This suggests that these two LLMs can become “confused” by too many ICL examples. This is reflected in Table 1 where there is a decrease in the accuracy value of Qwen on SemEval-2019 from $k = 10$ to $k = 20$.

Additionally, Appendix J highlights that the LLM indeed backs off to its parametric memory

when using low-quality demonstrations in ICL. This was validated by comparing the outputs of the LLM with poor-quality examples with a zero-shot configuration of the same LLM.

J Empirical Validation that LLMs under ICL with Poor-Quality Examples Back off to Parametric Memory

To validate the hypothesis that LLMs, when using ICL with poor-quality examples, back off to parametric memory, we compare their outputs with the same LLM in a zero-shot configuration. If our hypothesis is correct, both settings should produce similar labels, with a high Kappa agreement between them. In particular, we ran the three LLMs for 50 test samples under two settings: one with no examples (zero-shot) and the other with 10 poor examples. Table 15 shows the Cohen’s Kappa agreement between the two settings for different LLMs.

Datasets	Gpt-4o-mini	Llama-3.1-8B	QWEN-7B-Chat
AG News	0.86	0.57	0.45
News Category	0.98	0.43	0.63
AuTexTification	0.96	0.09	0.05
Implicit Prompts	0.69	0.59	0.40
SST-2	0.60	0.68	0.96
SemEval-2019	0.13	0.45	0.42

Table 15: Kappa agreement between zero-shot and poor-quality examples for all datasets.

The vast majority of Kappa values in Table 15 indicate substantial to near-perfect agreement, which aligns with our expectation that LLM relies on parametric memory when example quality is low, resulting in behavior comparable to the zero-shot. This is an encouraging result, which suggests that LLMs can ignore low-quality examples. There are some exceptions to the rule, which suggests that in some situations, low-quality demonstrations influence LLM behavior, like the SemEval-2019 dataset on Gpt-4o-mini, and AuTexTification on Llama-3.1-8B and QWEN-7B-Chat.

K Selected Examples Do not Necessarily Share the Same Label

A potential concern from our observations is that the similar behavior of LLM-based classifiers might be because of the possibility that selected examples in ICL or kNN¹⁵ have all the same label. However, as shown below, this largely does not apply to our scenario. To verify this, we analyze the proportion of test data points in each dataset whose

¹⁵We report results obtained using an unweighted kNN approach in this section.

selected examples have the same versus differing labels (see Table 16).

Datasets	$k = 10$	$k = 20$	$k = 30$
AG News	49.35	33.37	23.676
News Category	34	18.8	10.8
AuTexTification	4.6	0.2	0
Implicit Prompts	15.95	3.987	0.42
SST-2	26.49	11.353	5.73
SemEval-2019	5.2	0.6	0.2

Table 16: Percentage of datapoints where all the examples have the same label for increasing k .

In each of these cases, we further check the percentage of datapoints where the labels predicted by kNN and LLM align. For example, in Table 17, for the 49.35% of datapoints where AG News has all the same labels (from Table 16), kNN and Gpt-4o-mini make the same prediction 99.4% of the time. Similarly, as shown in Table 18, among the 50.65% of data points where AG News has different labels, kNN and Gpt-4o-mini agree in their predictions 81.85% of the time.

Datasets	$k = 10$	$k = 20$	$k = 30$
AG News	99.4	99.7	100
News Category	100	100	100
AuTexTification	84.78	100	0
Implicit Prompts	98.86	100	100
SST-2	87.45	94.95	96
SemEval-2019	69.23	83.33	100

Table 17: Percentage of datapoints out of the ones that have the *same* labels, where the prediction of kNN and Gpt-4o-mini match.

Datasets	$k = 10$	$k = 20$	$k = 30$
AG News	81.85	84.56	85.34
News Category	75.53	79.66	81.03
AuTexTification	54.61	53.61	55.4
Implicit Prompts	91.52	92.7	93.51
SST-2	65.83	69.47	70.92
SemEval-2019	47.58	49.8	46.5

Table 18: Percentage of datapoints out of the ones that have *different* labels, where the prediction of kNN and Gpt-4o-mini match.

While we expect the predictions to align well with the labels of examples in the first scenario, the second scenario also shows a high level of agreement. This supports our argument, as kNN and

Datasets	$k = 10$	$k = 20$	$k = 30$
AG News	100	100	100
News Category	100	100	100
AuTexTification	100	100	0
Implicit Prompts	95.83	98.48	100
SST-2	95.23	96.97	98.0
SemEval-2019	71.15	100	100

Table 19: Percentage of datapoints out of the ones that have the *same* labels, where the prediction of kNN and Llama-3.1-8B match.

Datasets	$k = 10$	$k = 20$	$k = 30$
AG News	83.83	87.26	88.09
News Category	84.44	86.64	86.83
AuTexTification	68.44	72.44	73.3
Implicit Prompts	78.15	83.57	83.43
SST-2	71.45	72.70	75.18
SemEval-2019	48.74	51.91	51.6

Table 20: Percentage of datapoints out of the ones that have the *different* labels, where the prediction of kNN and Llama-3.1-8B match.

Datasets	$k = 10$	$k = 20$	$k = 30$
AG News	100	100	100
News Category	81.87	86.17	93.52
AuTexTification	86.95	100	0
Implicit Prompts	35.60	27.27	14.28
SST-2	85.71	95.96	96
SemEval-2019	55.77	66.67	100

Table 21: Percentage of datapoints out of the ones that have the *same* labels, where the prediction of kNN and QWEN-7B-Chat match.

Datasets	$k = 10$	$k = 20$	$k = 30$
AG News	85.01	89.06	89.79
News Category	61.78	72.55	77.23
AuTexTification	59.54	59.32	61.4
Implicit Prompts	19.84	18.82	22.88
SST-2	64.27	68.43	70.19
SemEval-2019	43.68	45.38	45

Table 22: Percentage of datapoints out of the ones that have the *different* labels, where the prediction of kNN and QWEN-7B-Chat match.

Gpt-4o-mini tend to make similar predictions even when the example labels are not identical.

Similar to Table 17 and 18, Table 19 and Table 20 present results for Llama-3.1-8B under the same scenario. Interestingly, the percentage of data points with different labels of examples in Table 20

for the labels predicted by kNN and the Llama-3.1-8B shows high agreement when compared to Table 18. Further, we extend the analysis to QWEN-7B-Chat and present results in Table 21 and 22. In this case as well, we observe a similar behavior as both Gpt-4o-mini and Llama-3.1-8B, indicating that the underlying patterns remain consistent across different settings.