

How DDAIR you?

Disambiguated Data Augmentation for Intent Recognition

Galo Castillo-López Alexis Lombard Nasredine Semmar Gaël de Chalendar

Université Paris-Saclay, CEA, List, Palaiseau, France

{galo-daniel.castillolopez, alexis.lombard, nasredine.semmar, gael.de-chalendar}@cea.fr

Abstract

Large Language Models (LLMs) are effective for data augmentation in classification tasks like intent detection. In some cases, they inadvertently produce examples that are ambiguous with regard to untargeted classes. We present DDAIR (Disambiguated Data Augmentation for Intent Recognition) to mitigate this problem. We use Sentence Transformers to detect ambiguous class-guided augmented examples generated by LLMs for intent recognition in low-resource scenarios. We identify synthetic examples that are semantically more similar to another intent than to their target one. We also provide an iterative re-generation method to mitigate such ambiguities. Our findings show that sentence embeddings effectively help to (re)generate less ambiguous examples, and suggest promising potential to improve classification performance in scenarios where intents are loosely or broadly defined.

1 Introduction

Large Language Models have been proposed and effectively used for data augmentation in several NLP tasks in low-resource scenarios, including intent recognition (Wang et al., 2024; Si et al., 2025; Meguellati et al., 2025). A common approach to generate synthetic data from LLMs is through multiple LLM calls, where a single example is generated per call, or a set of examples belonging to a target class are generated in the same inference. While such a practice can produce useful synthetic utterances for intent recognition (Benayas et al., 2024), the lack of information about other intents in the generative process may cause ambiguous¹ outputs. For instance, an augmented utterance that is generated to represent a target intent, could be semantically more similar to another class. Arguably, such ambiguities negatively impact classification systems (Liu et al., 2023).

¹In this work, our definition of “ambiguous” is task-related rather than linguistic.

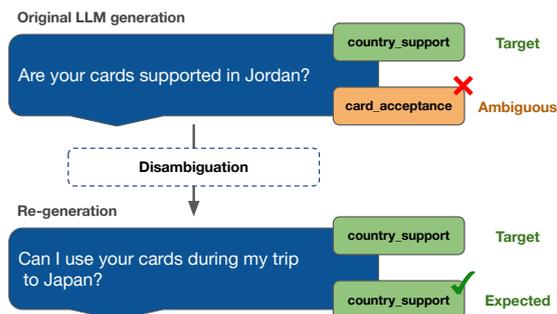


Figure 1: **Top:** A LLM-generated utterance for the target label `country_support` from the banking domain². The synthetic utterance is semantically similar to the `card_acceptance` intent. **Bottom:** After one disambiguation iteration, the re-generated utterance shows no ambiguity with respect to any class from the label space.

In this paper, we study the impact of ambiguous augmented examples from LLMs for intent recognition. We propose re-generating detected ambiguous utterances to improve augmented data quality. Our method, named DDAIR, consists of combining sentence transformers and LLMs to detect and re-generate ambiguous synthetic examples. We first employ sentence transformers to automatically detect ambiguous LLM-generated utterances, by semantically comparing synthetic utterances with few-shot examples from the training set (Figure 1 upper part). Then, detected ambiguous examples are re-generated using the LLM, including in-context learning examples, the target intent and the ambiguous (unexpected) intent names in the prompt (Figure 1 lower part). Our approach is iterative, such that the detect-and-generate process can be repeated as long as the augmented utterance remains ambiguous.

The main goal of our work is to investigate how detecting and treating ambiguous utterances impacts systems which are fine-tuned with examples generated by LLMs. We assess our method on

²Label examples from the BANKING77 corpus.

models from two LLM families, and three sentence encoders. Our results show that sentence transformers can be effectively used to detect ambiguous generations, and in combination with LLMs, reduce the ambiguity of augmented data. The implementation of the experiments conducted in this work is publicly available³.

2 Related Work

Recent work has made key progress on intent detection (Atuhurra et al., 2024). LLMs have been widely used as intent classifiers (Arora et al., 2024; Park et al., 2025; Castillo-López et al., 2025), whereas their application for data augmentation in low-resource scenarios remains relatively under-explored (Lin et al., 2024; Beatricia et al., 2024). Previous studies have found that LLMs consistently perform better than traditional methods at augmenting data for intent recognition (Benayas et al., 2024). Sahu et al. (2022) observed that the examples generated by LLMs are less helpful in tasks with semantically close intents, as models tend to generate utterances that belong to a semantically related intent to the target class.

Generating and refining synthetic examples for intent detection was studied in (Lin et al., 2024). Their two-stage approach involves using LLMs to generate utterances, followed by fine-tuning a smaller sequence-to-sequence model in a full-shot setting for unseen intents. In contrast, our work assumes limited data scenarios and does not rely on full-shot settings. Lin et al. (2023) propose using Pointwise \mathcal{V} -Information (Ethayarajh et al., 2022) to identify and retain the most informative synthetic utterances for training, leading to improved classification performance. While their approach filters out low-quality augmented data, our method focuses on enhancing it.

3 Experimental Procedure

We propose a data augmentation method for intent recognition in few-shot scenarios. The intent recognition task involves classifying a user utterance u into an intent label $y \in \mathcal{Y}$. To address data scarcity, we leverage large language models to generate synthetic utterances conditioned on intent names to form an augmented dataset. We propose the use of sentence transformers to encode both the synthetic utterances and the in-context learning examples to

detect ambiguous (i.e., problematic) utterances and provide a disambiguation strategy.

3.1 Utterance Generation with LLMs

We employ an in-context learning approach (ICL) to generate synthetic utterances on the instruct versions of two LLMs: Mistral 7B and Llama-3 8B. Our experiments consider from 2 to 5 randomly selected (without replacement) ICL examples per intent class from the training set to simulate a n -shot scenario. To reduce the variability introduced by random sampling, we carry out 5 independent sampling rounds for each ICL setting. This results in 5 distinct sets of ICL examples for each n -shot setup, where $n \in \{2, 3, 4, 5\}$, enabling a more robust evaluation. We perform one LLM call for each generated utterance. Generation prompts include the task instruction, ICL examples and intent name. More information about our use of LLMs is described in Appendix A.2.

3.2 Ambiguous Utterance Detection

In this work, we propose the detection and treatment of ambiguous generations to improve the quality of synthetic data. We use sentence transformers to encode the generated utterances and ICL examples, to detect ambiguous examples. We consider that a generated utterance is ambiguous if it lies far from the n ICL examples of the target intent (in some embedding space) and closer to the examples of an untargeted intent. We hypothesize that such ambiguous examples can damage the performance of intent recognition models. In particular, let $f(\cdot)$ denote the sentence encoder and $f(u)$ represent the embedding of an utterance u . For every utterance u_i in the training set, we obtain its representation $f(u_i)$. For each intent label $y \in \mathcal{Y} = \{1, \dots, m\}$, we compute a class-specific centroid vector c_y by averaging the embeddings of its corresponding ICL examples. Let $\{u_1^{(y)}, \dots, u_n^{(y)}\}$ denote the set of n ICL utterances labeled with intent y , then the centroid for class y is then computed as $c_y = \frac{1}{n} \sum_{i=1}^n f(u_i^{(y)})$. This results in one centroid vector $c_y \in \mathbb{R}^d$ for each intent class $y \in \mathcal{Y}$, where d is the sentence embedding dimensionality. We define a synthetic utterance u_j as ambiguous if $\hat{y}_j \neq y_j$, where $y_j \in \mathcal{Y}$ is its target intent and $\hat{y}_j \in \mathcal{Y}$ is its nearest centroid’s class, i.e., the label of the nearest centroid does not match its target label. We use BGE, MPNet and MiniLM-L6 to compute sentence embeddings. See A.3 for more details about the sentence transformers.

³<https://github.com/gaalocastillo/ddair>

3.3 Synthetic Utterance Disambiguation

We follow two strategies to treat ambiguous generations. First, similarly as in (Lin et al., 2023), we drop all ambiguous generated examples produced by the LLMs. Hence, we only keep synthetic utterances that lie in the neighborhood of the target intent ICL examples. The second approach consists in a multi-step strategy of detection and re-generations. We detect ambiguous generated utterances, as described in Section 3.2, and use a disambiguation –i.e., re-generation– prompt (see A.2). Our aim is to semantically align the synthetic utterances to the correct intent label. We conduct the disambiguation procedure iteratively up to 3 times in this work.

3.4 Intent Classification

To assess the utility of our proposed data augmentation approach, we fine-tune uncased BERT_{BASE} on the classification task (Devlin et al., 2019). Models are fine-tuned using both the few-shot and generated utterances. We also conduct a set of evaluations on baseline experiments using the original LLM generations –without any disambiguation strategy– with varying numbers of ICL examples and generated utterances, which can be found in Appendix B. See Appendix A.1 for more details about the fine-tuning process.

3.5 Datasets

Three intent detection corpora are used in this work: BANKING77 (Casaneva et al., 2020), CLINC150 (Larson et al., 2019) and the MPGT corpus (Addlessee et al., 2023). The BANKING77 corpus consists of 77 intents from the banking domain. CLINC150 comprises 150 intents from 10 domains, such as travel, work, and others. In addition, we use MPGT, a collection of multi-party dialogues between users and a hospital receptionist robot. The original MPGT corpus considers that utterances may have more than one label. Hence, we use the multi-class version of the corpus, proposed in (Castillo-López et al., 2025), which comprises 8 intents. More information about the datasets in Appendix F.

3.6 Evaluation

The evaluation of our work is three-fold. We first evaluate how our iterative disambiguation approach impacts the utterances’ ambiguity. To do so, we compute the ambiguity ratio as the proportion of ambiguous utterances from all the LLM-generated examples after each iteration. Second, we perform

a Silhouette coefficient analysis, a widely used clustering validation metric balancing intra-cluster cohesion and inter-cluster separation (Rousseeuw, 1987). The metric is regarded as a robust indicator of cluster quality (Vendramin et al., 2010; Arbelaitz et al., 2013) and has recently been used in intent discovery to assess embedding-space cluster structure (Liu et al., 2021; De Raedt et al., 2023; Hong et al., 2025; Ferrera et al., 2025)

Since clusters are given by the known intent labels, we do not perform clustering ourselves. Instead, we apply Silhouette analysis to assess whether the generated utterances align closely with the ICL examples of their expected intents or drift toward neighboring intent regions. For a given utterance i , let $a(i)$ be the mean distance between i and all other utterances within the same cluster (i.e., same intent), and let $b(i)$ be the minimum mean distance between i and all utterances in any other cluster (i.e., the nearest neighboring intent cluster). The Silhouette Coefficient $s(i)$ is defined as:

$$s(i) = \frac{b(i) - a(i)}{\max\{a(i), b(i)\}}$$

The coefficient ranges from -1 to 1 , where values close to 1 indicate that an utterance is well matched to its assigned cluster, values near 0 suggest a boundary case, and negative values imply potential misassignment.

Finally, we evaluate our fine-tuned BERT intent classifiers using the average macro-F1 over 3 runs with different seeds. As described in §3.1, we sample 5 random n -shot sets, so each reported score averages 15 results. We also implement and evaluate the data-augmentation method from (Lin et al., 2023) under our settings for comparison.

4 Results

Sections 4.1 and 4.2 provide two different analyses to assess our disambiguation strategy to reduce ambiguity on augmented utterances. We evaluate the utility of our disambiguation strategy to produce higher quality data for the intent classification task in Section 4.3.

4.1 Ambiguity Ratio

Figure 2 displays the ambiguity ratios after each disambiguation step. We observe that ambiguity ratios decrease after every disambiguation step in all scenarios. Our results suggest that we can use sentence transformers alongside LLMs to identify

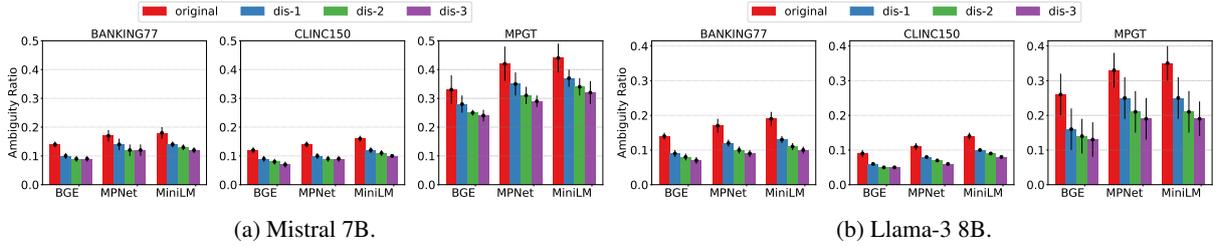


Figure 2: Ambiguity ratios of the original generations and re-generations after multiple iterative disambiguation steps (dis-1, dis-2, dis-3). A lower ratio indicates a lower proportion of ambiguous generated utterances. In all scenarios, ambiguity ratios decrease after disambiguation steps. MPGT is the most ambiguous corpus.

n -shot	Dis. Strategy	BANKING77			CLINC150			MPGT		
		BGE	MPNet	MiniLM	BGE	MPNet	MiniLM	BGE	MPNet	MiniLM
Mistral 7B										
2	w/o dis.	56.62 ± .40	56.62 ± .40	56.62 ± .40	75.82 ± .19	75.82 ± .19	75.82 ± .19	62.17 ± 1.92	62.17 ± 1.92	62.17 ± 1.92
2	PVI (*)	32.99 ± .49	32.99 ± .49	32.99 ± .49	46.11 ± .30	46.11 ± .30	46.11 ± .30	56.22 ± 1.52	56.22 ± 1.52	56.22 ± 1.52
2	Drop	51.12 ± .33	49.58 ± .44	48.42 ± .52	67.90 ± .37	67.33 ± .27	66.11 ± .16	68.00 ± 1.18	75.15 ± 1.07	69.75 ± 1.72
2	Dis-1	57.43 ± .38	57.03 ± .30	57.28 ± .37	75.95 ± .14	76.11 ± .14	76.27 ± .15	68.08 ± 1.14	68.19 ± 1.19	69.54 ± 1.44
2	Dis-2	57.37 ± .42	56.96 ± .31	56.93 ± .38	76.19 ± .16	76.10 ± .18	76.40 ± .18	68.36 ± .88	70.43 ± .93	69.67 ± 1.23
2	Dis-3	57.07 ± .44	56.66 ± .34	57.16 ± .27	76.43 ± .18	76.32 ± .13	76.56 ± .17	65.81 ± 1.45	72.05 ± 1.07	68.98 ± 1.21
5	w/o dis.	71.14 ± .17	71.14 ± .17	71.14 ± .17	85.11 ± .11	85.11 ± .11	85.11 ± .11	75.69 ± .49	75.69 ± .49	75.69 ± .49
5	PVI (*)	55.33 ± .37	55.33 ± .37	55.33 ± .37	67.72 ± .43	67.72 ± .43	67.72 ± .43	73.08 ± 0.79	73.08 ± 0.79	73.08 ± 0.79
5	Drop	70.45 ± .19	70.24 ± .23	70.51 ± .21	83.27 ± .14	82.29 ± .18	81.85 ± .13	81.27 ± .65	79.25 ± .46	78.38 ± .79
5	Dis-1	71.53 ± .23	72.12 ± .15	71.67 ± .18	85.09 ± .15	85.12 ± .11	85.11 ± .10	76.78 ± .52	76.67 ± .52	76.05 ± .73
5	Dis-2	71.53 ± .20	71.89 ± .18	71.64 ± .17	85.23 ± .10	85.17 ± .08	85.07 ± .12	77.28 ± .50	76.71 ± .64	77.03 ± .51
5	Dis-3	71.61 ± .18	71.70 ± .15	71.67 ± .16	85.31 ± .12	85.22 ± .12	85.16 ± .13	77.28 ± .50	75.40 ± .48	76.03 ± .83
Llama-3 8B										
2	w/o dis.	59.30 ± .26	59.30 ± .26	59.30 ± .26	77.19 ± .24	77.19 ± .24	77.19 ± .24	70.95 ± 1.09	70.95 ± 1.09	70.95 ± 1.09
2	PVI (*)	34.02 ± .54	34.02 ± .54	34.02 ± .54	44.45 ± .48	44.45 ± .48	44.45 ± .48	56.45 ± 1.30	56.45 ± 1.30	56.45 ± 1.30
2	Drop	52.90 ± .42	50.74 ± .34	49.53 ± .44	71.29 ± .28	69.75 ± .23	67.25 ± .22	72.73 ± .39	72.16 ± .11	71.94 ± .88
2	Dis-1	60.33 ± .28	59.98 ± .29	60.84 ± .31	77.67 ± .17	77.41 ± .21	78.24 ± .21	71.56 ± .96	71.50 ± .91	73.04 ± .89
2	Dis-2	60.31 ± .16	60.45 ± .27	60.35 ± .26	77.74 ± .21	77.51 ± .19	77.93 ± .19	69.17 ± 1.05	71.45 ± .81	70.86 ± .30
2	Dis-3	60.60 ± .15	60.35 ± .20	60.52 ± .23	77.95 ± .17	77.48 ± .15	78.13 ± .22	70.96 ± 1.00	71.78 ± .36	72.81 ± .51
5	w/o dis.	72.14 ± .19	72.14 ± .18	72.14 ± .18	85.26 ± .13	85.26 ± .13	85.26 ± .13	80.26 ± .47	80.26 ± .47	80.26 ± .47
5	PVI (*)	55.12 ± .41	55.12 ± .41	55.12 ± .41	66.54 ± .36	66.54 ± .36	66.54 ± .36	71.53 ± 0.98	71.53 ± 0.98	71.53 ± 0.98
5	Drop	70.38 ± .16	70.15 ± .16	69.65 ± .22	84.13 ± .16	83.81 ± .15	83.01 ± .14	83.04 ± .48	80.54 ± .98	82.76 ± .42
5	Dis-1	72.57 ± .18	72.30 ± .17	72.50 ± .17	85.36 ± .13	85.32 ± .15	85.32 ± .12	80.02 ± .47	80.36 ± .56	78.83 ± .52
5	Dis-2	72.64 ± .16	72.54 ± .13	72.54 ± .16	85.35 ± .15	85.33 ± .15	85.54 ± .14	79.73 ± .39	79.95 ± .50	79.05 ± .50
5	Dis-3	72.70 ± .12	72.72 ± .12	72.30 ± .16	85.48 ± .13	85.41 ± .14	85.50 ± .15	80.50 ± .47	80.27 ± .51	80.70 ± .28

Table 1: Results on the intent classification models, fine-tuned on n few-shot and 10 synthetic examples. Scores correspond to the average macro F1 and their standard deviations over 15 runs. Strategies include: not conducting any disambiguation step (w/o dis.); dropping ambiguous generations (Drop); multiple iterative disambiguation steps (Dis-1, Dis-2, Dis-3); and our implementation of the Per-Intent PVI method from (Lin et al., 2023) (*). Scores in **bold** highlight the best performance, and underlined scores represent the second best performance in the n -shot setup, using a given generative model and a sentence transformer for disambiguation. Results show that it is useful to process ambiguities: the two best performances are always obtained after dropping or disambiguation.

ambiguous synthetic examples and generate less ambiguous utterances for intent recognition. These findings are in line with our Silhouette coefficient analysis in Section 4.2. We identify that the MPGT corpus exhibits the highest ambiguity ratios in all settings, presenting ratios about twice as large as the other datasets in similar scenarios. For instance, more than 40% of the original Mistral generations on MPGT corresponds to ambiguous utterances according to MPNet and MiniLM-L6, while less than 20% of utterances are labeled as ambiguous on the other datasets. We also note that ambiguity ratios decrease as the number of ICL examples increases in all our settings of LLMs, datasets and embedding spaces. Figure 3 illustrates such behaviour on the original generations from Mistral 7B for the CLINC150 corpus.

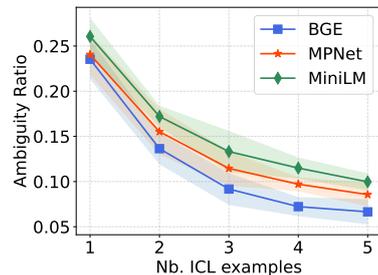


Figure 3: Ambiguity ratio by number of in-context learning examples on Mistral 7B generations for the CLINC150 corpus on multiple embedding spaces.

4.2 Silhouette Coefficient Analysis

Figure 4 shows that, under the same experimental settings, disambiguation generally improves cluster cohesion and separation, corroborating our ambiguity ratio findings. On BANKING77 and

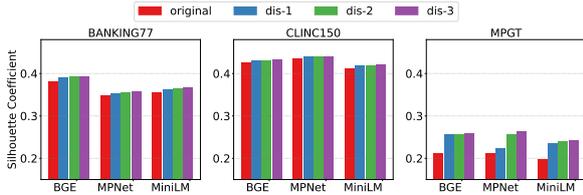


Figure 4: Silhouette coefficients of the original Llama-3 8B generations and re-generations after multiple iterative disambiguation steps (dis-1, dis-2, dis-3). A higher coefficient indicates a better inter-cluster and intra-cluster mean distance relation of the utterance in the embedding space. In all scenarios, the coefficients increase after 3 disambiguation steps.

CLINC150, gains are modest but consistent across encoders. On MPGT, where the baseline Silhouette coefficient is the lowest, improvements are substantial. Corresponding results for Mistral, reported in Figure 8 (Appendix C), show a first-step regression on MPGT across all encoders, with later steps recovering and surpassing the baseline for MPNet and BGE, while MiniLM-L6 remains close to the baseline. Overall, performing 3 disambiguation steps outperforms the no-disambiguation condition in 17 of 18 settings, indicating the benefit of iterative disambiguation, especially in domains with overlapping intents.

4.3 Intent Classification

Our results on the thorough set of experiments without disambiguation, reported on Appendix B, show that better performances are obtained when increasing the numbers of ICL and synthetic examples. Thus, we evaluate our proposed disambiguation strategies considering 2 and 5 ICL examples, as well as 10 augmented examples. Table 1 shows the evaluation of the intent recognition models in 2-shot and 5-shot scenarios. We observe that although performance gains are not large when following disambiguation steps in **5-shot** settings, we consistently obtain improvements on the BANKING77 and CLINC150 datasets after 3 steps for both generators. In most setups (11 out of 12), the best performing models are obtained after 2 or 3 disambiguation iterations in both corpora. On the other hand, our results show that the best strategy for the MPGT corpus is dropping ambiguous generations in all setups, whereas such a strategy is the worst for the other datasets. We obtain improvements on MPGT of up to ≈ 6 and ≈ 3 points on Mistral and Llama generations, respectively. We also find that our methods outperforms our implemen-

tation of the PVI-based approach in all setups. We believe that filtering out examples, through the PIV-based approach, results in class imbalance problems at fine-tuning. We provide further analysis on the PVI results in Appendix D.

Similar to our results in 5-shot, we see that the best scores in all the **2-shot** experiments are obtained after following one of the disambiguation strategies. For instance, models that are fine-tuned with Mistral generations on MPGT present significant enhancements after dropping or re-generating utterances. Average macro-F1 is increased to up to ≈ 6 points on MPGT. Moreover, we note that larger standard deviations are found on 2-shot experiments than the 5-shot scenario, which might be explained by the instability of the intent centroids given the lower number of examples. We hypothesize that higher performance lifts are seen in the MPGT corpus, compared to the other corpora, due to the high ambiguity on its synthetic examples, as well as to the presence of lower granularity on intent definitions than the other corpora (Sahu et al., 2022). This hypothesis is supported by our silhouette coefficient analysis, which shows the largest improvements on MPGT where baseline clustering quality was poorest. We provide further analysis on MPGT ambiguities in Appendix E. Additionally, qualitative evaluation of our proposed method is provided in Appendix G.

5 Conclusions

We present an iterative data augmentation method to detect and re-generate ambiguous utterances for the intent recognition task. Our work shows that sentence encoders, in combination with LLMs, can be used to reduce the ambiguity ratio of generated utterances and produce examples with better quality for intent detection systems. In line with previous research, we observe that our method significantly improves the performance of systems with loosely or broadly defined intents. We also acknowledge that our proposed disambiguation approach may result in a higher number of LLMs calls. Nevertheless, we argue that such additional computational costs in the generative task might be amortized in the long term. We discuss such a trade-off in Appendix H. We hope that our findings provide insights for future work in data augmentation, not only for intent recognition, but also other text classification tasks.

Limitations

Ambiguity Definition. As explained in Section 3.2, we define an utterance as *ambiguous* whether its closest intent centroid, in the selected embedding space, is not its target intent. One important limitation on that definition is the sensibility to outliers of the intent centroid computation. Particularly, there might be cases where certain few-shot examples lie far from the rest of the examples used to compute the centroid. This may cause instability, especially in highly data-scarce scenarios, such as 1-shot or 2-shot settings. We aim to mitigate this issue by computing the element-wise median of the sentence embeddings as an alternative to obtain a more robust estimate of the central tendency for each intent class, notably in the presence of outliers. Let $\{u_1^{(y)}, \dots, u_n^{(y)}\}$ denote the set of n ICL utterances labeled with intent y , and let $f(u_i^{(y)})$ denote their sentence embeddings. We define the class-specific median vector \tilde{c}_y as the element-wise median across the embeddings. This results in one median vector $\tilde{c}_y \in \mathbb{R}^d$ for each intent class $y \in \mathcal{Y}$, where the median is taken independently over each of the d embedding dimensions. By analyzing the computed medians, we observed that our results, based on the centroids, were not sensible to outliers. Nonetheless, we believe that future work may consider different strategies to detect ambiguous utterances on the embedding space, such as k-NN, or an adapted k-NN algorithm on the centroids – the k-Nearest Centroids.

Dependency on Embeddings Quality. Our proposed ambiguity detection method heavily relies on the sentence embedding space. However, we aim to provide conclusive findings, regardless of the selected embedding space, by conducting experiments over several well-known encoders combined with various generators.

Ethical Considerations

Our work essentially consists in refining text generation from Large Language Models for data augmentation. Synthetic utterances generated by LLMs, as other types of LLM-generated texts, may present inherent biases from the models. Hence, outputs from our proposed method rely in controlled generation strategies to mitigate bias and harmful LLM outputs. After a manual analysis of the generated text in our experiments, we do not foresee harm resulting from the experiments

conducted in this study. Additionally, our experiments use publicly available corpora, which have been curated prior to our work to prevent malicious actions.

The LLM inference performed in this work to generate synthetic utterances was executed on private infrastructure using a single NVIDIA A100 SXM 40 GB GPU. The infrastructure has a carbon efficiency of 0.432 kgCO₂eq/kWh. The total time required for the generative tasks across all experiments was 23.6 hours using Mistral 7B and 32.3 hours using Llama-3 8B. Therefore, the total emissions are estimated at 6.03 kgCO₂eq. These estimations are based on the Machine Learning Impact calculator⁴ (Lacoste et al., 2019).

Acknowledgments

We warmly thank our anonymous reviewers for their time and valuable feedback. This publication was made possible by the use of the FactoryIA supercomputer, financially supported by the Ile-de-France Regional Council. This work has been partially funded by the EU project CORTEX² (under grant agreement: N° 101070192). Some mathematical expressions included in Section 3 were written with the assistance of ChatGPT⁵.

References

- Angus Addlesee, Weronika Sieińska, Nancie Gunson, Daniel Hernandez Garcia, Christian Dondrup, and Oliver Lemon. 2023. [Multi-party goal tracking with LLMs: Comparing pre-training, fine-tuning, and prompt engineering](#). In *Proceedings of the 24th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 229–241, Prague, Czechia. Association for Computational Linguistics.
- Olatz Arbelaitz, Ibai Gurrutxaga, Javier Muguerza, Jesús M Pérez, and Iñigo Perona. 2013. An extensive comparative study of cluster validity indices. *Pattern recognition*, 46(1):243–256.
- Gaurav Arora, Shreya Jain, and Srujana Merugu. 2024. [Intent detection in the age of LLMs](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 1559–1570, Miami, Florida, US. Association for Computational Linguistics.
- Jesse Atuhurra, Hidetaka Kamigaito, Taro Watanabe, and Eric Nichols. 2024. Domain adaptation in intent classification systems: a review. *arXiv preprint arXiv:2404.14415*.

⁴<https://mlco2.github.io/impact/>

⁵<https://chatgpt.com/>

- Sherleen Beatricia, Albert Reinaldo, Muhammad Fikri Hasani, and 1 others. 2024. Text augmentation for Indonesian intent classification: Comparative study. In *2024 11th International Conference on Information Technology, Computer, and Electrical Engineering (ICITACEE)*, pages 265–270. IEEE.
- Alberto Benayas, Sicilia Miguel-Ángel, and Marçal Mora-Cantallops. 2024. Enhancing intent classifier training with large language model-generated data. *Applied Artificial Intelligence*, 38(1):2414483.
- Iñigo Casanueva, Tadas Temčinas, Daniela Gerz, Matthew Henderson, and Ivan Vulić. 2020. [Efficient intent detection with dual sentence encoders](#). In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*, pages 38–45, Online. Association for Computational Linguistics.
- Galo Castillo-López, Gael de Chalendar, and Nasredine Semmar. 2025. [Intent recognition and out-of-scope detection using LLMs in multi-party conversations](#). In *Proceedings of the 26th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 504–512, Avignon, France. Association for Computational Linguistics.
- Maarten De Raedt, Frédéric Godin, Thomas Demeester, and Chris Develder. 2023. [IDAS: Intent discovery with abstractive summarization](#). In *Proceedings of the 5th Workshop on NLP for Conversational AI (NLP4ConvAI 2023)*, pages 71–88. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pages 4171–4186.
- Kawin Ethayarajh, Yejin Choi, and Swabha Swayamdipta. 2022. [Understanding dataset difficulty with \$\mathcal{V}\$ -usable information](#). In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 5988–6008. PMLR.
- Alessandra Ferrera, Giulio Mezzotero, and Domenico Ursino. 2025. [A linguistics-based approach to refining automatic intent detection in conversational agent design](#). *Information Sciences*, 689:121493.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Mengze Hong, Wailing Ng, Chen Jason Zhang, Yuanfeng Song, and Di Jiang. 2025. [Dial-In LLM: Human-Aligned LLM-in-the-loop Intent Clustering for Customer Service Dialogues](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 5885–5900, Suzhou, China. Association for Computational Linguistics.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *Preprint*, arXiv:2310.06825.
- Alexandre Lacoste, Alexandra Luccioni, Victor Schmidt, and Thomas Dandres. 2019. Quantifying the carbon emissions of machine learning. *arXiv preprint arXiv:1910.09700*.
- 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, Hong Kong, China. Association for Computational Linguistics.
- I-Fan Lin, Faegheh Hasibi, and Suzan Verberne. 2024. [Generate then refine: Data augmentation for zero-shot intent detection](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 13138–13146, Miami, Florida, USA. Association for Computational Linguistics.
- Yen-Ting Lin, Alexandros Papangelis, Seokhwan Kim, Sungjin Lee, Devamanyu Hazarika, Mahdi Namazifar, Di Jin, Yang Liu, and Dilek Hakkani-Tur. 2023. [Selective in-context data augmentation for intent detection using pointwise V-information](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1463–1476, Dubrovnik, Croatia. Association for Computational Linguistics.
- Alisa Liu, Zhaofeng Wu, Julian Michael, Alane Suhr, Peter West, Alexander Koller, Swabha Swayamdipta, Noah Smith, and Yejin Choi. 2023. [We’re afraid language models aren’t modeling ambiguity](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 790–807, Singapore. Association for Computational Linguistics.
- Pengfei Liu, Youzhang Ning, King Keung Wu, Kun Li, and Helen Meng. 2021. Open intent discovery through unsupervised semantic clustering and dependency parsing. *IEEE CogInfoCom-2021*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.

- Elyas Meguellati, Assaad Zeghina, Shazia Sadiq, and Gianluca Demartini. 2025. Llm-based semantic augmentation for harmful content detection. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 19, pages 1190–1209.
- Gyutae Park, Ingeol Baek, Byeongjeong Kim, Joongbo Shin, and Hwanhee Lee. 2025. [Dynamic label name refinement for few-shot dialogue intent classification](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 41–52, Vienna, Austria. Association for Computational Linguistics.
- 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.
- Peter J Rousseeuw. 1987. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20:53–65.
- Gaurav Sahu, Pau Rodriguez, Issam Laradji, Parmida Atighehchian, David Vazquez, and Dzmitry Bahdanau. 2022. [Data augmentation for intent classification with off-the-shelf large language models](#). In *Proceedings of the 4th Workshop on NLP for Conversational AI*, pages 47–57, Dublin, Ireland. Association for Computational Linguistics.
- Yusuf Sali and Sitki Can Toraman. 2025. [Navigating the unknown: Intent classification and out-of-distribution detection using large language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 14652–14664, Suzhou, China. Association for Computational Linguistics.
- Lijia Si, Caili Guo, Zheng Li, and Yang Yang. 2025. A unified framework of data augmentation using large language models for text-based cross-modal retrieval. *Pattern Recognition*, page 111755.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2020. MpNet: Masked and permuted pre-training for language understanding. *Advances in neural information processing systems*, 33:16857–16867.
- Lucas Vendramin, Ricardo JGB Campello, and Eduardo R Hruschka. 2010. Relative clustering validity criteria: A comparative overview. *Statistical analysis and data mining: the ASA data science journal*, 3(4):209–235.
- Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. 2020. Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. *Advances in neural information processing systems*, 33:5776–5788.
- Yusong Wang, Dongyuan Li, Jialun Shen, Yicheng Xu, Mingkun Xu, Kotaro Funakoshi, and Manabu Okumura. 2024. Lambda: large language model-based data augmentation for multi-modal machine translation. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 15240–15253.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and 1 others. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 conference on empirical methods in natural language processing: system demonstrations*, pages 38–45.
- Shitao Xiao, Zheng Liu, Peitian Zhang, Niklas Muenninghoff, Defu Lian, and Jian-Yun Nie. 2024. [C-Pack: Packed Resources For General Chinese Embeddings](#). In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '24*, pages 641–649, New York, NY, USA. Association for Computing Machinery.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, and 1 others. 2022. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.

A Model Information

In this appendix we provide model implementation details of our experiments.

A.1 BERT

We use BERT to build intent recognition models. All our runs of BERT fine-tuning were conducted on a single NVIDIA Tesla V100 SXM2 GPU of 32GB. Uncased BERT_{BASE} (110M parameters) is used in all the experiments. We used the BertForSequenceClassification class from Hugging Face’s Transformers library (Wolf et al., 2020) for sequence classification tasks. Table 2 shows the hyperparameter configuration we employ. For all intent classification experiments, we report the results averaged over three runs, using random seeds ranging from 0 to 2.

A.2 Large Language Models

We use Large Language Models to generate – and re-generate– synthetic utterances for data augmentation. Our experiments use instruct versions of Mistral 7B and Llama-3 8B. Specifically, we use Mistral-7B-Instruct-v0.2 (Jiang et al., 2023) and Meta-Llama-3-8B-Instruct (Grattafiori et al., 2024). Figures 5 and 6 show the prompt templates we use on the generation and re-generation experiments, respectively.

hyperparameter	value
eval_monitor	macro F1
train_batch_size	16
eval_batch_size	16
test_batch_size	16
wait_patience	3
num_train_epochs	25
warmup_proportion	0.1
lr	1e-5

Table 2: Set of hyperparameters used on BERT fine-tuning experiments.

Generation Prompt

Give me 1 user’s utterance indicating the user intent of "[intent]" in the domain "[domain]". You can use the following utterances as examples:

- ICL example 1
- ICL example 2
- ...
- ICL example N

Your response should only be a JSON object with the following structure:

```
{"utterance": "generated_utterance"}
```

Figure 5: Prompt template used on all LLM generation experiments. Highlighted text in blue varies according to the expected intent and in-context learning examples.

A.3 Sentence Transformers

Our proposed methods consist in detecting ambiguous augmented examples for intent recognition. To do so, we use sentence embeddings to encode both real and synthetic utterances, and capture semantic representations of training examples. Experiments are performed with three sentence transformers: BAAI/bge-base-en-v1.5 (Xiao et al., 2024), all-mpnet-base-v2 (Song et al., 2020), and all-MiniLM-L6-v2 (Wang et al., 2020). We used the SentenceTransformer class from the SentenceTransformers library (Reimers and Gurevych, 2019) to compute utterance embeddings.

B Baseline Experiments

We perform an exhaustive evaluation on the use of multiple ICL examples and generated utterances,

Re-generation Prompt

I generated 1 user’s utterance indicating the user intent of "[intent]" in the domain "[domain]" as follows: "[ambiguous_utterance]". However, the utterance is more similar to the intent "[most_similar_intent]" which may cause ambiguities. Please refine and disambiguate the generated utterance. You can use the following utterances as examples of the expected intent:

- ICL example 1
- ICL example 2
- ...
- ICL example N

Your response should only be a JSON object with the following structure:

```
{"utterance": "generated_utterance"}
```

Figure 6: Prompt template used on all LLM re-generation experiments (disambiguation). Highlighted text in blue varies according to the expected intent, the computed most similar intent, the ambiguous synthetic utterance, and in-context learning examples.

without any disambiguation strategy. In particular, we evaluate classification models on different combinations of ICL examples –from zero to five– and synthetic utterances –from zero to ten. Figure 7 indicates that, consistently across all datasets, higher amounts of in-context learning instances and generated utterances, provide better performances.

C Silhouette Coefficient Analysis

Figure 8 shows how the silhouette coefficient varies after each disambiguation step. We note that the coefficient tends to increase as more disambiguation steps are conducted.

D PVI Experiments

D.1 Implementation Details

We compare our proposed approach with the method proposed in (Lin et al., 2023). To do so, we conduct our own implementation of the Pointwise \mathcal{V} -Information (PVI) method (Ethayarajh et al., 2022). PVI is a score that computes the amount of information that is carried by an instance for a classification task. In other words, the score can be

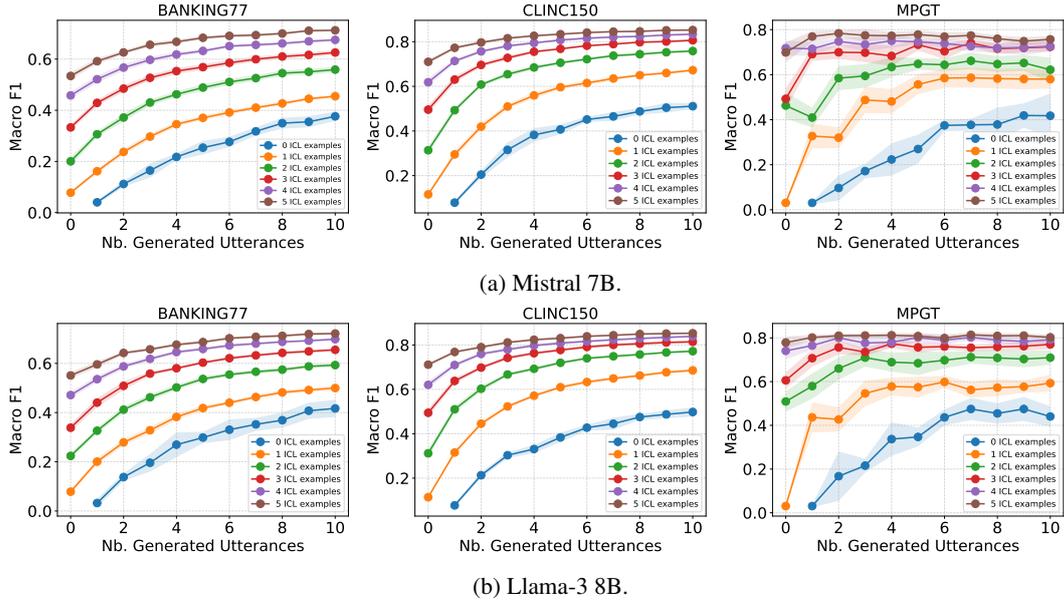


Figure 7: Macro-F1 score of the intent recognition models fine-tuned on the generated utterances. We evaluate different combinations on the number of augmented examples (from zero to ten) and amount of ICL examples (from zero-shot to five-shot settings). Increasing both numbers provides higher performances.

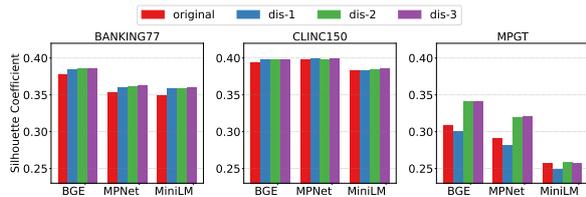


Figure 8: Silhouette coefficients of the original Mistral 7B generations and re-generations after multiple iterative disambiguation steps (dis-1, dis-2, dis-3). A higher coefficient indicates a better inter-cluster and intra-cluster mean distance relation of the utterance in the embedding space. In all scenarios, the coefficients increase after 3 disambiguation steps.

used to quantify how useful a synthetic utterance is for an intent recognition task. A higher value indicates that the example provides more information, and thus that it is more useful. We use the same models described in A.1 and A.2 to compute the PVI on generated utterances. Lin et al. (2023) proposed two PVI methods: Global PVI, which uses a threshold based on all the examples; and Per-Intent PVI, which uses a different threshold per class. In Table 1, we report the results obtained using Per-Intent PVI. We invite readers to refer to (Lin et al., 2023; Ethayarajh et al., 2022), for more details about the PVI-based method.

D.2 Results Analysis

Intent detection models’ evaluations reported in Table 1, show that our approach significantly outperform our implementation of the PVI-based method. We argue that filtering out certain synthetic utterances (e.g. less informative examples), may result in problems related to class imbalance at training time. We note that, by following the PVI-based method on a given corpus, the number of examples per intent varies between 2 and 9, out of 10 originally generated utterances. Figure 9 shows the distribution of unfiltered synthetic examples per class on both the PVI-Intent and Drop methods. We observe that the PVI-Intent approach exhibits an under-representation of more classes than our Drop strategy. In other words, certain classes are relatively more represented than others at fine-tuning in the classification task. These findings are observed in all corpora, across all LLM and sentence encoder settings. Note that our re-generation strategies do not discard examples, contrary to the PVI-based filtering and Drop methods. Thus, the re-generation strategy do not induce any class imbalance.

Additionally, we believe that the computed performances are lower than the ones reported in (Lin et al., 2023) because our experiments consider different configurations. For instance, their implementation use larger models than the ones used in this work: OPT-66B (Zhang et al., 2022) and RoBERTa_{LARGE} (Liu et al., 2019) as generator and

classification model, respectively. We have made every effort to faithfully implement the method proposed in Lin et al. (2023). However, as with any reproduction, there remains the possibility of unintentional discrepancies.

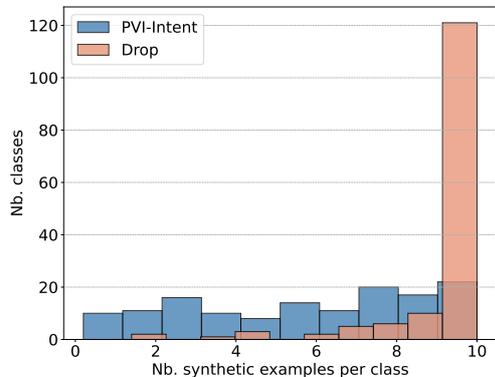
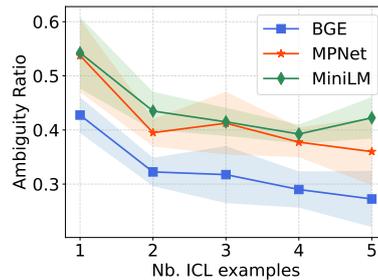


Figure 9: Distribution of unfiltered examples per class after applying the PVI-Intent and Drop (with respect to BGE) methods on Mistral 7B generations over the CLINC150 corpus. The PVI-Intent method under-represents more intents than the Drop approach.

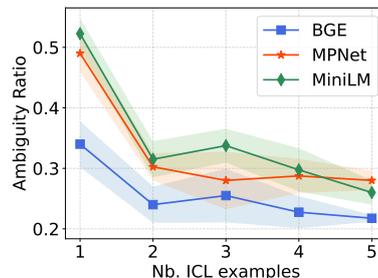
E MPGT Results Analysis

We observed in Table 1 that the MPGT corpus shows the highest performance lifts when conducting a disambiguation strategy. One key difference between the BANKING77 or CLINC150 datasets, and MPGT, is the semantic similarity among intents. While BANKING77 and CLINC150 comprise 77 and 150 intents, respectively, the MPGT dataset includes only 8 intents. Intents in the MPGT corpus include: affirm, bye thank, cant-help, incomplete-da, inform, offer-help, request, suggest. On the other hand, BANKING77 corresponds to one single domain, and CLINC150 to 10. We argue that our method shows performance gains in scenarios with loosely or broadly defined intents.

We also found that the MPGT corpus presents a higher ambiguity ratio, in comparison to the other corpora. Figure 10 shows the ambiguity ratio by number of ICL examples. Following (Ethayarajh et al., 2022), we compute the Pointwise \mathcal{V} -Information (PVI) score to assess the information carried by the synthetic examples for each corpus for our classification task. Figure 11 indicates that the synthetic examples for the MPGT dataset present the lowest PVI scores, evidencing that generating useful utterances for the intent recognition task is more difficult on the MPGT corpus. Also,



(a) Mistral 7B.



(b) Llama-3 8B.

Figure 10: Ambiguity ratio by number of in-context learning examples on LLM generations for the MPGT corpus on multiple embedding spaces. A lower ratio indicates a lower proportion of ambiguous generated utterances.

we see that the PVI scores on the 5-shot setups are much larger than in the 2-shot scenarios, which confirms that using more ICL examples lets the generators produce more useful utterances.

F Dataset Statistics

Table 3 shows the number of examples per dataset split.

	train	dev	test
BANKING77	8K	2K	3K
CLINC150	15K	3K	5K
MPGT	.5K	.1K	.2K

Table 3: Number of examples per dataset split.

G Re-generation Qualitative Analysis

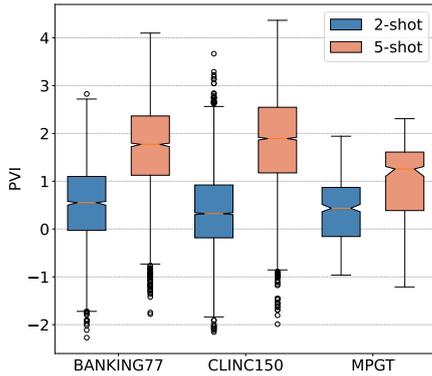
Table 4 shows examples of ambiguous utterances disambiguated in one and two disambiguation steps. The table includes the original utterance generation (iteration 1), the target label, the most similar label with respect to BGE, and the subsequent re-generation iterations. We observe that ambiguities occur for two main reasons: the use of ambiguous terms that may relate to unintended labels, and the

inherent semantic overlap of intents in the label space.

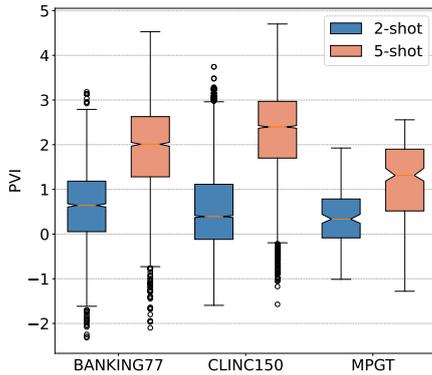
H Disambiguation Computational Cost

In this section we discuss the cost-benefit trade-off of our disambiguation strategy. While our approach increases the number of LLM calls for data augmentation, the motivation of this work is to produce higher-quality training data that enables small, efficient models –e.g., BERT-based classifiers– to achieve substantially better intent recognition performance without relying on LLMs at inference time, as proposed in previous work (Arora et al., 2024; Castillo-López et al., 2025; Sali and Toraman, 2025). Therefore, the one-time augmentation cost is amortized over the lifetime of a much more efficient deployment, an advantage aligned with prior work showing the long-term utility of high-quality synthetic data for small, or hybrid (LLMs + smaller models), intent recognition systems.

In addition, only a subset of examples require more than one iteration, and empirically this fraction is small in corpora with large label spaces such as BANKING77 and CLINC150. Thus, the overhead is limited in practice. The computational cost due to our re-generative approach is directly proportional to the ambiguity ratio described in Figure 2. Table 5 shows the cumulative percentual additional costs due to the iterative procedure with respect to the original data augmentation, when using Mistral 7B and BGE on all the datasets.



(a) Mistral 7B.



(b) Llama-3 8B.

Figure 11: PVI scores of the synthetic examples for each dataset in 2-shot and 5-shot settings. Higher values indicate that utterances are more informative (i.e., more useful) for the classification task. The MPGT corpus shows the lowest PVI values, suggesting a higher difficulty to generate good quality utterances.

	BANKING77	CLINC150	MPGT
dis-1	13.2%	11.6%	31.3%
dis-2	23.2%	19.9%	58.1%
dis-3	32.2%	27.3%	81.7%

Table 5: Computational costs of the disambiguation strategy with Mistral 7B and BGE on all the datasets. Values correspond to the cumulative percentual additional costs across iterations (costs are accumulated after each iteration) due to the iterative procedure **with respect to the original data augmentation**, i.e. without disambiguation.

dataset	original (it. 1)	target label	sim. label (it. 1)	re-generation (it. 2)	sim. label (it. 2)	re-generation (it. 3)	sim. label (it. 3)
BANKING77	What places accept	card	supported cards	Where can I use	card	–	–
	payments with my banking card?	acceptance	and currencies	this banking card for payments?	acceptance		
CLINCI150	What are the home-related tasks that I need to remember to do?	reminder	todo list reminder	What are the home-related tasks that need to be reminded to be done?	reminder	–	–
CLINCI150	What is the due date for my next bank account payment?	bill due	payday	What is the due date for my next bank account payment?	payday	What is the due date for my next bank account bill payment?	bill due
CLINCI150	Can you show me my spending history on transactions in the last month?	spending history	transactions	What are the details of my spending history on transactions in the last month in the bank?	transactions	What are the details of my spending in the last month, in food and transportation?	spending history
MPGT	Yes, I can help with that.	affirm	offer help	Absolutely, I can help you with that.	offer help	Absolutely	affirm

Table 4: Qualitative analysis of our disambiguation strategy. Original utterances and re-generations correspond to the first generations (iteration 1) and subsequent disambiguated utterances (iterations 2 and 3) produced by Mistral 7B, respectively. The target label is included, as well as the most similar label computed after each iteration with respect to BGE.