



HAL
open science

Investigating Low-Cost LLM Annotation for Spoken Dialogue Understanding Datasets

Lucas Druart, Valentin Vielzeuf, Yannick Estève

► **To cite this version:**

Lucas Druart, Valentin Vielzeuf, Yannick Estève. Investigating Low-Cost LLM Annotation for Spoken Dialogue Understanding Datasets. 27th International Conference on Text, Speech and Dialogue, Sep 2024, Brno (Rép. Tchèque), Czech Republic. hal-04616105

HAL Id: hal-04616105

<https://hal.science/hal-04616105v1>

Submitted on 19 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Investigating Low-Cost LLM Annotation for Spoken Dialogue Understanding Datasets

Lucas Druart^{1,2}, Valentin Vielzeuf²[0000-0002-2987-2510]
Yannick Estève¹[0000-0002-3656-8883]

¹ Laboratoire d’Informatique d’Avignon (LIA), Avignon, France
`first.last@univ-avignon.fr`

² Orange Innovation, Rennes, France
`first.last@orange.com`

Abstract. In spoken Task-Oriented Dialogue (TOD) systems, the choice of the semantic representation describing the users’ requests is key to a smooth interaction. Indeed, the system uses this representation to reason over a database and its domain knowledge to choose its next action. The dialogue course thus depends on the information provided by this semantic representation. While textual datasets provide fine-grained semantic representations, spoken dialogue datasets fall behind. This paper provides insights into automatic enhancement of spoken dialogue datasets’ semantic representations. Our contributions are three fold: (1) assess the relevance of Large Language Model fine-tuning, (2) evaluate the knowledge captured by the produced annotations and (3) highlight semi-automatic annotation implications.

Keywords: spoken dialogue systems, automatic annotation, large language models, spoken language understanding

1 Introduction

Digitization enables many tasks to be automated, nevertheless users sometimes require assistance to perform complex tasks such as making a reservation at a restaurant or booking a hotel room. Task-Oriented Dialogue (TOD) systems are designed to assist such users. A common approach to implement them is to break the problem down to three iterative steps [23]: updating the system’s understanding of the users’ needs, reasoning over a database and domain knowledge to choose the next action and providing the user an answer. Those systems often rely on transfer learning which requires annotated datasets. However only few datasets provide aligned dialogue recordings with turn-level contextual semantic annotations. Therefore the dialogue understanding community has mainly focused on textual datasets creating a usage discrepancy [8].

Dialogue systems rely on a chosen semantic representation to infer the next action(s). The gap between textual semantic representations and spoken ones provides an explanation for the observed discrepancy. Indeed, spoken dialogue

datasets often use a flat semantic representation embedded in the transcription, such as transcription span labels, while textual ones propose more fine-grained structured representations such as Dialogue-Abstract Meaning Representation [3] or Dialogue Meaning Representation [13]. This paper proposes a method to automatically annotate dialogue datasets with fine-grained semantic representations.

```
<command-tache> j' aurais voulu réserver > euh <nombre-chambre-reservation>
deux > <chambre-type> chambres doubles > <connectProp> et >
<nombre-chambre-reservation> une > <chambre-type> chambre simple >
<sejour-nbNuit-reservation> pour les cinq jours >
<temps-jourFerie-reservation> de Noël à la Noël > donc
<localisation-arrondissement-hotel> dans le huitième arrondissement >
<localisation-ville-hotel> de Paris >
```

(a) Current labeled span annotation.

```
(r1 / reservation
  :objet (h1 / hotel
    :chambre (e1 / et
      :arg1 (c1 / chambre
        :type "double"
        :quantite "deux")
      :arg2 (c2 / chambre
        :type "simple"
        :quantite "une"))
    :date-sejour (e2 / evenement
      :nom "Noël")
    :duree-sejour (d1 / duree
      :quantite "cinq"
      :unite "jours")
    :lieu (a1 / adresse
      :ville "Paris"
      :quartier "huitième arrondissement"))
  :etat "en cours")
```

(b) Targeted semantic tree annotation.

Fig. 1: Example of enriched annotation for a user turn of the MEDIA dataset. It can be translated as "I would like to book two double bedrooms and one single bedroom for the five days of Christmas at Christmas so in the eighth district of Paris.". Node **identifiers** are in red, **node types** in blue, **relation types** in green and structure in black. Transcription spans are quoted.

Some recent datasets attempt to foster work on spoken dialogues: A spoken version of the MultiWOZ dataset [4] with vocalized user turns was published in the context of the *Speech Aware Dialogue Systems* track of the 11th edition of

the Dialogue System Technology Challenge³ (DSTC11) [21]. SpokenWOZ [20] provides 5,700 dialogue recordings annotated with the same annotation format as MultiWOZ *i.e.* list of slot-value pairs. However none are as fine-grained as Abstract Meaning Representation (AMR) [1]. The STOP dataset [22] provides structured semantic trees for single turn voice commands which unfortunately leads to only 16% truly requiring trees (*i.e.* depth greater than 2) [6]. Conversely, spoken TOD are well suited for structured semantic trees since sub-tasks are discussed in sub-dialogues which complement each other.

The high cost of fine-grained annotations seems to be the hurdle preventing spoken dialogue datasets from embracing such formalism. Yet, unannotated data is much more affordable and semi-supervised approaches gather a large panel of tasks [16,25,10], leading to a low-cost compromise towards automatic annotation. With the rise of large generative models, automatic annotation of textual data is thus reaching a new level of possibilities. Indeed, finetuning Large Language Models (LLM) on consumer-grade GPUs [12] is now a reality and has been applied to complex tasks such as Machine Translation [18]. Moreover, prompt-based approaches [19,26,11] even offer zero-shot annotations. Nevertheless, when dealing with complex tasks, fully automatic annotation remains challenging, because of behaviors such as hallucinations [24] and recent work often keeps a human in the loop [15] or go back to former semi-supervised methods.

2 Method

Given the heavy cognitive load required to annotate dialogue turns with a contextual version of AMR, we attempt to automatize as much as we can while keeping high quality expectations. To do so we define a structured contextual meaning representation fitting the dataset’s use case in section 2.1, set an annotation pipeline in section 2.2 and evaluate this system in section 3.2.

2.1 Structured Contextual Meaning Representation

Two adaptations are required to fit our dataset’s use case: defining an ontology fitting the hotel booking task and handling cross-turn references.

The defined ontology comprises three type of concepts: domain related concepts (*e.g.* hotel) which represent hotel reservation elements, operators (*e.g.* et) which enable to apply an operation to other concept(s) and general purpose concepts (*e.g.* adresse) which are domain-agnostic. Each concept has a unique identifier which enables exact cross-turn references. Concepts can be linked together or to transcription spans called literals. The edge’s label depends on the pair of concepts considered.

Compared with the current labeled span annotation, this annotation takes into account the interactions between concepts and highlights the implicit concepts which are not uttered but group other concepts together (*e.g.* hotel in

³ <https://dstc11.dstc.community/>

Figure 1). Those implicit concepts have higher error rates [17] and become essential to understand short turns with indirect references such as "the one with a swimming pool".

2.2 Annotation Pipeline

Our annotation pipeline comprises the following four steps which can be repeated over several iterations as illustrated in Figure 2.




Fig. 2: Overview of the semi-automatic annotation pipeline with the parameters of each step.

Human annotation. The first step consists in having 10 trained human annotators annotate a subset of the corpus following the defined ontology. We focused on the dataset’s test set which comprises 208 dialogues with 30% of them annotated by multiple annotators. The reported annotator agreement reaches an average semantic match (smatch) [5] of 77.28%. We present a few statistics of the obtained annotations in Table 1 which highlight the importance of structured annotation. For each dialogue, we select the annotation with the least errors, according to a master-annotator, as the final one. Our goal is to leverage this set of clean annotations to produce the remaining annotations.

In order to track the quality of the automatic annotations, we further annotate another set of 22 dialogues to form a 10% fold of unseen annotations.

Avg. # user turns	width > 2	depth > 2
16.95	31.86%	24.34%

Table 1: Annotated dialogue turns statistics.

Fine-Tuning. We then fine-tune a Mistral-7B LLM [14] on the human annotated data with Low Rank Adaptation (LoRA) [12]. Given a pre-trained weight matrix $\mathbf{W} \in \mathbb{R}^{d \times k}$, we fine-tune a delta matrix $\Delta\mathbf{W} = \mathbf{B}\mathbf{A}$ such that $\mathbf{B} \in \mathbb{R}^{d \times r}$ and $\mathbf{A} \in \mathbb{R}^{r \times k}$ with $r \ll \min(d, k)$. We then only train this factorized matrix $\Delta\mathbf{W}$ with the forward pass thus modified to:

$$\mathbf{y} = \mathbf{W}\mathbf{x} + \frac{\alpha}{r}\Delta\mathbf{W}\mathbf{x}, \text{ with } \alpha = mr, m \in \mathbb{N}$$

The LLM is fed the following prompt template in which it is tasked to provide a structured annotation of the last turn of a sequence of t (**agent**, **user**) speaker turns transcription pairs⁴.

Prompt Template

Below is an instruction that describes a task, paired with an input that provides further context. Write a response that appropriately completes the request.

Instruction: Provide the tree annotation of what is said by the user in the given dialogue.

Input: agent: [agent₀]; user: [user₀] ... agent: [agent_t]; user: [user_t]

Response:

Constrained Generation. In order to reduce hallucinations and ensure that every output is correctly formatted while preserving originality, we implement a grammar constrained decoding [9] and allow up to 256 newly generated tokens. At each decoding step, the grammar constraint simply sets the log-probabilities of the forbidden tokens of the vocabulary to $-\infty$. The next token’s probability distribution is thus restricted to grammar valid tokens. The main challenge consists in decoding smaller units (*i.e.* tokens) than the grammar terminals while ensuring that the decoding remains in valid grammar states, especially when reaching literals which are supposed to be open vocabulary. In that case, our constraint decoding only allows tokens from the speakers turns until another quote is decoded. While more restrictive, this ensures that literals match speaker’s transcription spans. Algorithm 1 presents how we select the set of allowed token level decoding paths to transition from one grammar valid state to another.

⁴ In practice the model is only fed the 5 previous turns to limit the number of tokens.

Algorithm 1: Ensures that the decoding remains in valid grammar states.

Data: W next grammar potential terminals, T model’s tokenizer, eos end-of-speech token

Result: map allowed tokens mapping for all decoding paths which reach another valid grammar state

```

 $map \leftarrow$  empty map;
if  $W$  is empty then
  |  $map[T(eos)] \leftarrow$  empty map;
else
  for  $w \in W$  do
    |  $tokens \leftarrow T(w)$ ;
    |  $next\_tokens \leftarrow map$ ;
    | while  $len(tokens) > 1$  do
      |  $t_0 \leftarrow tokens[0]$ ;
      | if  $t_0 \notin next\_tokens$  then
        | |  $next\_tokens[t_0] \leftarrow$  empty map;
        | |  $tokens \leftarrow tokens[1 : ]$ ;
        | |  $next\_tokens \leftarrow next\_tokens[t_0]$ ;
      |  $next\_tokens[tokens[0]] \leftarrow$  empty map;

```

Annotation Evaluation. Finally we evaluate the generated annotations with the AMR Semantic Match (smatch) score [5]. This metric searches for the best variable alignment and then computes the F1 score of the matching triples for this alignment. While it accounts for potential variable permutations, it remains a matching metric. Indeed, an error in the name of a concept or in a literal value invalidates the whole triple. Yet, averaged over many annotations, it provides a valuable insight in the structural quality of a set of annotations. It thus helps us to choose the best LLM annotator and to quantify the annotation quality over the held-out dialogues. Additionally, we track the number of ontology errors.

3 Results

3.1 Dataset

The MEDIA dataset [7] comprises over 1,000 French hotel reservation dialogue recordings annotated with turn-level span concept labels over the manual transcriptions. It is recognized as a challenging spoken dialogue understanding dataset [2]. Indeed the Wizard-of-Oz collection protocol provided loosely scripted scenarios which fostered rich and natural interactions.

3.2 Experimental Setup

This paper addresses three concerns around the annotation pipeline described in the sections above:

1. Is LLM fine-tuning relevant for this setup?
2. Is the knowledge captured by the LLM effectively transferred to the automatic annotations?
3. What does this mean for semi-automatic annotation?

To assess the relevance of LLM fine-tuning, we compute pairwise annotation comparisons produced by a human A, a human B and a LLM X which is either fine-tuned or prompted in a few-shot manner.

To evaluate the knowledge captured by the automatic annotations, we fine-tune the backbone LLM over the automatic annotations and compare its quality with the one which provided the annotations.

Finally, to grasp the implications for semi-automatic annotation, we train a score estimator and use it to filter out the worst annotations (*i.e.* score lower than the threshold δ) before iterating over the model training and annotation. The estimator is composed of a Support Vector Regressor (SVR) over sentence embeddings⁵ of the last dialogue turn and the produced annotation.

3.3 Relevance of LLM Fine-Tuning

We first compare a standard prompting strategy to fine-tuning one. Prompting can be done on any model including commercial ones which are often among the top performers of leaderboards. However, their token pricing policies may become prohibitive for iterative pipelines. On the other hand, fine-tuning requires a GPU and careful hyper-parameters setting but enables fine-grained customization.

Figure 3 presents the similarity (computed with smatch) distributions of pairwise comparisons between human and automatic annotations. We observe that a fine-tuned Mistral-7B can achieve higher similarities than Chat-GPT 3.5. However both LLMs remain far from human-wise annotation similarities indicating that iterative annotation should be promising. Therefore we choose the fine-tuning approach over the prompt-based one.

3.4 Knowledge captured by automatic annotations

We perform a hyper-parameter grid search over $\alpha \in [r, 2r]$, $r \in [16, 128, 512]$ and learning rate $\eta \in [1, 4, 8] \times 10^{-4}$ and select $r = \alpha = 512$ and $\eta = 4 \times 10^{-4}$. We present the results obtained with this model over the clean set and the unseen set in Table 2.

We first observe the importance of the dialogue history for such fine-grained annotations as the **Clean last turn** model which is only provided the last dialogue turn performs worse than the **Clean history** model.

Then, the grammar constrained decoding seems to behave as a complex temperature: it tends to make the model more verbose thus never predicting empty trees. This constrained decoding seems to disturb the calibration of the model’s conditional probabilities. It improves the full tree annotations but makes the model more verbose hence reducing the performance on empty trees.

⁵ Obtained with the model at <https://huggingface.co/thenlper/gte-large>




Fig. 3: Smatch scores distribution of pairwise comparisons of human and automatic annotations.

Training set	Clean		Unseen 10%	
	Full	Empty	Full	Empty
Clean last turn	52.43 +/-18.82	94.78 +/-9.94	33.12 +/-12.57	90.06 +/-11.73
<i>w. grammar</i>	73.31 +/-10.84	0.0	60.94 +/-10.02	0.0
Clean history	86.85 +/-10.06	92.99 +/-11.31	58.44 +/-13.29	78.2 +/-23.83
<i>w. grammar</i>	82.17 +/-10.92	0.0	66.66 +/-7.62	0.0
Iteration 1	49.93 +/-15.93	86.56 +/-15.13	43.65 +/-14.66	81.69 +/-17.06
<i>w. grammar</i>	64.36 +/-11.95	0.0	60.99 +/-13.0	0.0
Merged	52.22 +/-15.09	86.56 +/-15.13	46.52 +/-14.31	81.69 +/-17.06

Table 2: Smatch scores for the same backbone model trained over different training sets and evaluated on the clean annotations and the unseen 10% fold.

Finally, the annotations produced by the **Clean history** model are of high enough quality to train well performing model. Indeed, when filtering out the worst examples as described above, the **Iteration 1** model comes close to its parent. Further iterations and more selective filtering therefore seem promising to improve the models' performances. In addition, the grammar constrained decoding seems more relevant with such noisier models. Merging both uncon-

strained and constrained predictions helps reach a happy-medium between full and empty trees as shown by the **Merged** model.

4 Discussion

This work focuses on the MEDIA dataset [7] since it is recognized as challenging [2] and hence provides complex enough situations to require fine-grained annotations. However it is a single domain French dataset. Performances might be better on higher resources languages such as English and worse on lower resources languages. While this paper focuses on ontology specific to MEDIA’s use case, designing a similar one for multi-domain use-cases should be manageable in most cases.

With the recent hype around Large Language Models, a huge quantity of resources has been made available. This paper only focuses on well-established models and techniques. To the best of our knowledge prompt design remains an unsettled topic. We thus experimented with several formulations and selected the one which seemed the most promising. There might be better formulations than the one proposed in this paper.

5 Conclusion

Our journey towards fine-grained annotations for spoken dialogue datasets has led us to the realm of LLMs in order to comply with our low annotation budget. This exploratory work provides valuable insights to the community on the design of complex automatic annotation for spoken dialogue datasets. Indeed, our method may accelerate manual annotation and/or be included in a fully automated setting by carefully selecting the training examples. We highlight that open-weights LLMs fine-tuning is relevant for this annotation setup since it enables faster iterations than prompting commercial models while remaining competitive. We also propose a grammar constrained decoding strategy which struggles with non-informative dialogue turns but improves the annotation of correctly annotated dialogue turns. Our produced annotations contain part of the knowledge from the human annotations, and can be effectively learned by a model.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Banarescu, L., Bonial, C., Cai, S., Georgescu, M., Griffitt, K., Hermjakob, U., Knight, K., Koehn, P., Palmer, M., Schneider, N.: Abstract meaning representation for sembanking. In: Proceedings of the 7th linguistic annotation workshop and interoperability with discourse. pp. 178–186 (2013)

2. Béchet, F., Raymond, C.: Benchmarking benchmarks: introducing new automatic indicators for benchmarking Spoken Language Understanding corpora. In: *Inter-Speech (2019)*
3. Bonial, C., Donatelli, L., Abrams, M., Lukin, S.M., Tratz, S., Marge, M., Artstein, R., Traum, D., Voss, C.: Dialogue-AMR: Abstract Meaning Representation for dialogue. In: *Proceedings of the Twelfth Language Resources and Evaluation Conference*. Marseille, France (2020)
4. Budzianowski, P., Wen, T.H., Tseng, B.H., Casanueva, I., Ultes, S., Ramadan, O., Gasic, M.: Multiwoz - a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling. In: *Conference on Empirical Methods in Natural Language Processing (EMNLP) (2018)*, <https://api.semanticscholar.org/CorpusID:52897360>
5. Cai, S., Knight, K.: Smatch: an evaluation metric for semantic feature structures. In: *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. pp. 748–752. Association for Computational Linguistics, Sofia, Bulgaria (Aug 2013), <https://www.aclweb.org/anthology/P13-2131>
6. Chen, X., Ghoshal, A., Mehdad, Y., Zettlemoyer, L., Gupta, S.: Low-resource domain adaptation for compositional task-oriented semantic parsing. In: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics (2020)
7. Devillers, L., Bonneau-Maynard, H., Rosset, S., Paroubek, P., McTait, K., Mostefa, D., Choukri, K., Charnay, L., Bousquet-Vernhettes, C., Vigouroux, N., Béchet, F., Romary, L., Antoine, J.Y., Villaneau, J., Vergnes, M., Goulian, J.: The french media/evalda project: the evaluation of the understanding capability of spoken language dialogue systems. In: *International Conference on Language Resources and Evaluation (2004)*
8. Faruqui, M., Hakkani-Tür, D.: Revisiting the boundary between ASR and NLU in the age of conversational dialog systems. *Computational Linguistics (2022)*
9. Geng, S., Josifoski, M., Peyrard, M., West, R.: Grammar-constrained decoding for structured NLP tasks without finetuning. In: *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics (2023)
10. Georgila, K., Lemon, O., Henderson, J., Moore, J.D.: Automatic annotation of context and speech acts for dialogue corpora. *Natural Language Engineering* **15**(3), 315–353 (2009)
11. Gong, T., Belanich, J., Somandepalli, K., Nagrani, A., Eoff, B., Jou, B.: Lanser: Language-model supported speech emotion recognition. *arXiv preprint arXiv:2309.03978 (2023)*
12. Hu, E.J., yelong shen, Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: LoRA: Low-rank adaptation of large language models. In: *International Conference on Learning Representations (2022)*, <https://openreview.net/forum?id=nZeVKeeFYf9>
13. Hu, X., Dai, J., Yan, H., Zhang, Y., Guo, Q., Qiu, X., Zhang, Z.: Dialogue meaning representation for task-oriented dialogue systems. In: *Findings of the Association for Computational Linguistics: EMNLP 2022*. Association for Computational Linguistics (2022)
14. Jiang, A.Q., Sablayrolles, A., Mensch, A., Bamford, C., Chaplot, D.S., de las Casas, D., Bressand, F., Lengyel, G., Lample, G., Saulnier, L., Lavaud, L.R., Lachaux, M.A., Stock, P., Scao, T.L., Lavril, T., Wang, T., Lacroix, T., Sayed, W.E.: Mistral 7b (2023)

15. Kim, H., Mitra, K., Chen, R.L., Rahman, S., Zhang, D.: Meganno+: A human-llm collaborative annotation system. arXiv preprint arXiv:2402.18050 (2024)
16. Liao, X., Zhao, Z.: Unsupervised approaches for textual semantic annotation, a survey. *ACM Computing Surveys (CSUR)* **52**(4), 1–45 (2019)
17. Mdhaffar, S., Pelloin, V., Caubrière, A., Laperrière, G., Ghannay, S., Jabaian, B., Camelin, N., Estève, Y.: Impact analysis of the use of speech and language models pretrained by self-supervision for spoken language understanding. In: *International Conference on Language Resources and Evaluation (LREC)* (2022)
18. Moslem, Y., Haque, R., Way, A.: Fine-tuning large language models for adaptive machine translation. arXiv preprint arXiv:2312.12740 (2023)
19. Savelka, J.: Unlocking practical applications in legal domain: Evaluation of gpt for zero-shot semantic annotation of legal texts. In: *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*. pp. 447–451 (2023)
20. Si, S., Ma, W., Gao, H., Wu, Y., Lin, T.E., Dai, Y., Li, H., Yan, R., Huang, F., Li, Y.: SpokenWOZ: A Large-Scale Speech-Text Benchmark for Spoken Task-Oriented Dialogue Agents. In: *NeurIPS Datasets and Benchmarks Track* (2023), <https://openreview.net/forum?id=viktK3nO5b>
21. Soltau, H., Shafran, L., Wang, M., Rastogi, A., Han, W., Cao, Y.: DSTC-11: Speech aware task-oriented dialog modeling track. In: *Proceedings of The Eleventh Dialog System Technology Challenge*. Association for Computational Linguistics (2023)
22. Tomasello, P., Shrivastava, A., Lazar, D., Hsu, P.C., Le, D., Sagar, A., Elkahky, A., Copet, J., Hsu, W.N., Adi, Y., Algayres, R., Nguyen, T.A., Dupoux, E., Zettlemoyer, L., Mohamed, A.: Stop: A dataset for spoken task oriented semantic parsing. In: *2022 IEEE Spoken Language Technology Workshop (SLT)* (2022). <https://doi.org/10.1109/SLT54892.2023.10022703>
23. Tur, G., De Mori, R.: *SLU in Commercial and Research Spoken Dialogue Systems*. Wiley Telecom (2011). <https://doi.org/10.1002/9781119992691.ch7>
24. Xu, Z., Jain, S., Kankanhalli, M.: Hallucination is inevitable: An innate limitation of large language models. arXiv preprint arXiv:2401.11817 (2024)
25. Yang, X., Song, Z., King, I., Xu, Z.: A survey on deep semi-supervised learning. *IEEE Transactions on Knowledge and Data Engineering* **35**(9), 8934–8954 (2022)
26. Zhang, Y., Yang, R., Xu, X., Xiao, J., Shen, J., Han, J.: Teleclass: Taxonomy enrichment and llm-enhanced hierarchical text classification with minimal supervision. arXiv preprint arXiv:2403.00165 (2024)