



HAL
open science

Increasing faithfulness in human-human dialog summarization with Spoken Language Understanding tasks

Eunice Akani, Benoît Favre, Frédéric Bechet, Romain Gemignani

► To cite this version:

Eunice Akani, Benoît Favre, Frédéric Bechet, Romain Gemignani. Increasing faithfulness in human-human dialog summarization with Spoken Language Understanding tasks. 2024. hal-04699645

HAL Id: hal-04699645

<https://hal.science/hal-04699645v1>

Preprint submitted on 17 Sep 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

INCREASING FAITHFULNESS IN HUMAN-HUMAN DIALOG SUMMARIZATION WITH SPOKEN LANGUAGE UNDERSTANDING TASKS

Eunice Akani^{1,2}, **Benoit Favre**¹, **Frederic Bechet**^{1,3}, **Romain Gemignani**²

¹ Aix-Marseille Univ, CNRS, LIS, Marseille, France

² Enedis, Marseille, France

³ International Laboratory on Learning Systems - ILLS - IRL CNRS, Montréal, Canada

eunice.akani@univ-amu.fr, firstname.lastname@lis-lab.fr, romain.gemignani@enedis.fr

ABSTRACT

Dialogue summarization aims to provide a concise and coherent summary of conversations between multiple speakers. While recent advancements in language models have enhanced this process, summarizing dialogues accurately and faithfully remains challenging due to the need to understand speaker interactions and capture relevant information. Indeed, abstractive models used for dialog summarization may generate summaries that contain inconsistencies. We suggest using the semantic information proposed for performing Spoken Language Understanding (SLU) in human-machine dialogue systems for goal-oriented human-human dialogues to obtain a more semantically faithful summary regarding the task. This study introduces three key contributions: First, we propose an exploration of how incorporating task-related information can enhance the summarization process, leading to more semantically accurate summaries. Then, we introduce a new evaluation criterion based on task semantics. Finally, we propose a new dataset version with increased annotated data standardized for research on task-oriented dialogue summarization. The study evaluates these methods using the DECODA corpus, a collection of French spoken dialogues from a call center. Results show that integrating models with task-related information improves summary accuracy, even with varying word error rates.

1 Introduction

Advancements in spoken language processing, made possible by low word-error-rate automatic speech transcription and generative language models, have enhanced the quality of outputs across various tasks, including dialog summarization. Dialog summarization involves generating a concise, abstract rendition of the dialogue exchanged among speakers in a multi-party spoken interaction. Due to discrepancies between transcript and summary styles, inaccuracies in the transcript, verbal hesitations by participants, and the generative process itself, dialog summarization is susceptible to "hallucinations", wherein generative models produce text containing nonfactual information when compared to the source.

A study by Cao et al. [2018] found that 30% of summaries generated by text summarization systems contained incorrect information, known as "hallucinations" Maynez et al. [2020]. Approaches to assessing summary faithfulness include textual entailment Falke et al. [2019], Maynez et al. [2020], entity analysis Nan et al. [2021], Ji et al. [2023], and question-answer verification Durmus et al. [2020]. Dialog summarization is growing in popularity but is challenged by dialog structure and diverse data types, such as customer service and technical discussions. Methods using auxiliary information like dialogue acts Goo and Chen [2018] or domain terminology Koay et al. [2020] have been proposed. There are fewer studies on hallucinations in dialog summarization compared to text. Wang et al. [2022] found that 35% of SAMSum dataset summaries were inconsistent with the source dialogues. Tang et al. [2022] identified eight types of factual errors in dialogue summarization, while Wang et al. [2022] identified six types and used a summary model based on conditional generation probabilities to evaluate the model's faithfulness.

In industrial applications, this issue may be more prevalent due to the use of smaller language models for cost, energy efficiency, and data protection. Leveraging domain knowledge can help evaluate faithfulness and improve summary accuracy.

In this paper, we aim to enhance dialog summarization faithfulness by integrating task-specific elements like caller intent and domain-specific named entities. Our study focuses on call centers, using the DECODA corpus Bechet et al. [2012], a dataset of goal-oriented human interactions with semantic annotations. To our knowledge, no similar large, natural English speech dataset exists for this research. We also propose a new dataset using data augmentation via LLM to augment the dataset and improve the performance of dialog summarization model.

2 Integration of task-related information in spoken dialog summarization

On goal-oriented human-human conversations, we can leverage semantic information proposed for performing SLU in human-machine dialogue systems for tasks such as transportation reservations [Hemphill et al., 1990] or restaurants/hotels [Devillers et al., 2004]. Typically, here, three semantic levels are defined: domain (the semantic context of the dialogue, such as flight bookings), intent (the nature of the request made in the dialogue, like confirmation or information inquiry), and entity/value pairs (the semantic relationships within intents, such as a destination in an itinerary request). As we focus on goal-oriented, we put ourselves in a customer service conversation. Thus, the call type can be seen as an intent in SLU and the domain named entities as a slot. We hypothesize that by incorporating task-specific semantic information into the summarization process, we can effectively address inconsistencies related to this information in the data. In this section, we will focus on the call type and entity level. Using call type and named entities is an excellent choice for our study as it is easily obtained and leveraging named entities provides a practical framework for our methodology. Nevertheless, other semantic information, such as semantic frames can be used.

2.1 Summarization based on task information

Identifying different call types and concept labels is essential for analyzing conversation recordings or transcripts, but this information often needs to be inferred. This can be done by implementing a separate prediction system on the transcript and integrating call types into the input data of the summarization system. The method is **Pipeline_C**: Predict semantic information C from dialog D such as $C = \text{intent}(D)$, then condition summary generation on C : $S = \text{summary}(D, \text{semantic-information}(D))$

2.2 Task-related information as a selection criteria

We propose using semantic information to select the most reliable summary from different summaries obtained by tweaking decoding parameters such as top-p. The method is based on call type prediction and the hallucination risk on entities.

Call type Based on the hypothesis that the generated summary should have the same call type as the dialog transcription, we suggest using call type as a summary selection criterion. Our method involves using a text classifier to predict the call type of the generated summary and comparing it to the predicted call type for the entire transcript. We handle multiple call types and uncertainty in call type classification by calculating the divergence between the probability distribution on all call types for both summary and dialog classifiers using the KL divergence [Kullback and Leibler, 1951]. For G and R the probability distribution given by the call type classifier on the generated summary and the entire dialog respectively, CT the set of call types; the KL divergence between G and R is defined as follows:

$$D_{KL}(G \parallel R) = \sum_{x \in CT} G(x) \log \left(\frac{G(x)}{R(x)} \right) \tag{1}$$

This allows us to choose the summary that minimizes the D_{KL} among the generated summaries.

Named entities We propose using the NEHR criterion [Akani et al., 2023], which measures the proportion of named entities not present in the source document, as a selection criterion for creating faithful summaries. This approach has shown favorable results in reducing hallucinated entities in text summarization. Thus, we decided to apply the same experiment to dialog summarization.

Combining NEHR and D_{KL} This study aims to create a summary that accurately captures the task’s semantic information. We hypothesize that combining criteria like NEHR and D_{KL} will yield a summary consistent with

the source document. For H the set of summaries sampled from the model, m the minimum NEHR value, $V = \{x \in H | NEHR(x) = m\}$ the set of summaries with the minimum NEHR, $D_{KL}(x \parallel R)$ the KL divergence between a summary x associated with the dialog transcription R , the final system will be.

$$\hat{s} = \min_{x \in V} D_{KL}(x \parallel R) \quad (2)$$

2.3 Task-related information as an evaluation criteria

We propose a metric CT-Acc, which measures the accuracy of the call type classifier used on the generated summary compared to the reference call type. The assumption is that the summary is more coherent with the call type classifier if the accuracy is higher.

We propose using a metric based on named entities to assess the coherence of generated summaries compared to reference summaries. The metric measures entity-level accuracy and aims to ensure that both summaries contain the same entities. We measure NE-P, NE-R, and NE-F1 metrics to represent precision, recall, and F1 between entities in the generated and reference summaries. Our approach differs from previous work Nan et al. [2021] as we include all types of entities, such as dates and numerals, which we believe are crucial in some scenarios.

3 Experiments

3.1 The DECODA Corpus

DECODA-v1-hum [Bechet et al., 2012] contains spoken conversations between agents of the Paris Transport Authority and users of Paris buses and metro lines. Each conversation has a manual text transcription and a short summary. The corpus covers various call types such as *Itinerary*, includes entities from a domain ontology such as *Transport*, and *Schedule*, and consists of three parts (DECODA-1, DECODA-2, DECODA-3) with annotated synopses available only in DECODA-1 and DECODA-3. Since DECODA-1 summaries are written in a non-literary style and with little detail, we decided to keep only those in DECODA-3 for system development. From the 500 dialogues of the DECODA-3, we kept 200 dialogues as an evaluation set (*test*), 200 for fine-tuning the system (*Hum.* for *human* annotated data) and 100 as a validation set to adjust the parameters. As not all DECODA summaries are written in a literary style and are sometimes non-existent for certain conversations, we create a new dataset with more annotated data for task-oriented dialog summarization.

DECODA-v2-aug We leverage prompt-based LLM to create summaries for the remaining DECODA conversations. The methodology entails processing the manual transcription with an instruction-based Language Model (LM) in a few-shot learning mode, providing the system with a few examples of target summaries and instructing it to produce similar summaries. To do so, we decide to use ChatGPT-3.5 from OpenAI¹ as the model have good capacity to produce coherent and fluent French text and perform various tasks without extensive finetuning. In addition to the gold synopses from previous train dataset (*Hum.*), we add the summaries produced by ChatGPT-3.5 on the manual transcriptions of DECODA-1 and DECODA-2. We add some pair (dialog-automatic summary) to the development dataset and keep the same test set to form the new dataset call *Aug.* for *augmented*. Appendix B shows the prompt used to generate summaries.

DECODA-v2-auto Then, we put ourselves in a scenario where manual transcription was unavailable and opted to use an ASR system for transcribing the dialog. We used WhisperX [Bain et al., 2023], a state-of-the-art model for ASR based on Whisper [Radford et al., 2022]. Whisper is an ASR system trained on a large and diverse dataset, enabling transcription in multiple languages and translation into English. WhisperX adds three components to Whisper. It is a practical choice for ASR in contexts where manual transcription is challenging to obtain. We applied WhisperX on the raw audio of the corpus *Aug.* to have the automatic transcription and called the corpus *Auto.* for *Automatic transcription*.

Table 1 shows statistics for the training and test sets. Synopses in the test set are generally longer than in the training set. This disparity arises because the augmented data does not consistently mimic the style of human summaries in DECODA, given its 1-shot approach. The different datasets will be made available for research purposes on task-oriented dialog and particularly on task-oriented dialog summarization.

¹<https://openai.com/index/chatgpt/>

	Hum.	Aug.	Test
#dialog	200	1390	200
#conv.len	545	470	496
#sum.lm	55.3	47.9	52.7

Table 1: DECODA distribution in the train (Hum. and Aug.) and the test set. #dialog: number of dialogs; #conv.len and #sum.len average conversation and summary words respectively.

3.2 Summarization and classification models

Automatic Summarization We trained the summarization systems using BARThez Kamal Eddine et al. [2021], a sequence-to-sequence model pre-trained on various French corpora introduced for automatic text summarization. It is a transformer-based model built on BART architecture Vaswani et al. [2017], Lewis et al. [2020]. It consists of 6 encoder and decoder layers. We used the pre-trained model provided by the authors available on the Hugging Face library² to train the models.

Call type classification Two classifiers based on CamemBERT-base Martin et al. [2020] were trained to classify the call-type: one taking as input automatic dialog transcripts (Conv-CT-classifier) and the other one, generated synopses (Syn-CT-classifier). Conv-CT-classifier is also used to predict the call type for the Pipeline_C model.

3.3 Data augmentation and ASR evaluation

We compute ROUGE, BERTScore, and the semantic score of text summarization systems trained on the different corpus versions to validate our data augmentation strategy using LLMs and ASR systems. We report all results in table 2. For automatic transcription, we report result for two different sizes of WhisperX (Tiny-T and Large-L) and ChatGPT-3.5 one short on WhisperX Large transcription. The findings indicate that data augmentation enhances summary scores while semantic scores remain constant, with a slight increase in NE-F1. In terms of automatic transcription, several metrics declined. However, the scores of the large model continue to be competitive, suggesting that the ASR system could be a viable option. ChatGPT-3.5 achieves the highest semantic scores in both manual and automatic transcriptions. This is unsurprising, considering the model’s numerous parameters and the extensive training data it has received. In next section, only DECODA-v2-auto with WhisperX large transcription will be used for experiments.

Transcription	Data	ROUGE-L	BERTScore	CT-Acc	NE-F1
Manual transcription	Hum.	23.6	33.1	0.77	0.49
	Aug.	29.1	38.9	0.77	0.51
	one-shot*	28.9	37.3	0.80	0.51
Automatic transcription	Auto. (T)	24.5	34.4	0.75	0.31
	Auto. (L)	27.7	37.1	0.76	0.43
	one-shot*	26.6	35.4	0.79	0.46

Table 2: Automatic evaluation on the test partition with Barthez fine-tuned on the different datasets (Hum., Aug. Auto.). Performance of chatGPT3.5 with *. WER score of WhisperX Tiny (T) and Large (L) are respectively 76.4, 40.0.

3.4 Impact of selection criterion semantic criteria

We trained Pipeline_C, the model that conditions the generation of the summary by the call type concatenated to the dialog using a separator. From this model, we generated various summaries using sampling decoding strategies and select the one according to the selection metrics presented in section 2.2. We report automatic results in table 3 and denote as min D_{KL} for summaries selected using D_{KL} and min NEHR for summaries selected using NEHR and min NEHR + D_{KL} for the combined criterion. We report the result for Pipeline_C model and the baseline model without selection. For Pipeline_C, we have similar scores while CT-Acc vary. We can see that D_{KL} -based selection increases CT-Acc, while NEHR-based selection increases both values. The combined criteria seems to have the same result than D_{KL} -based selection.

²<https://huggingface.co/moussaKam/barthez>

System	ROUGE-L	BERTScore	CT-Acc	NE-F1
Baseline	27.7	37.1	0.76	0.43
Pipeline _C	27.6	37.0	0.79	0.43
min NEHR	27.7	36.4	0.81	0.46
min D_{KL}	27.8	36.8	0.82	0.44
min NEHR + D_{KL}	27.9	37.0	0.82	0.44

Table 3: BARThez + WhisperX Large + Pipeline_C - Summaries evaluation of summary selection method.

4 Discussion and Conclusion

We analyzed the influence of task-semantic information on the generation of dialog summaries. Our findings revealed that incorporating NEHR in this model enhances named entity precision. Nonetheless, manual evaluation is necessary to validate the effectiveness of this selection criterion in terms of fidelity and informativeness. In addition, we proposed new version of a dataset with more annotated data for task-oriented dialog summarization. Evaluation of the dataset revealed that using LLM to generate a portion of training synopses improved the model’s performance, as indicated by increased ROUGE score and BERTScore, and automatic transcriptions can serve as a viable alternative to address the lack of annotated data.

Acknowledgements

This work was granted access to the HPC resources of IDRIS under the allocation 2023-AD011012525R3 made by GENCI.

A Details about DECODA dataset

	DECODA-1	DECODA-2	DECODA-3
Style	Synthetic	×	Literary
#dialog	1009	478	500
#turn	54.6	71.5	78.1
#conv.len	411.8	513.3	504.1
#sum.len	24.7	×	53.8

Table 4: Description of DECODA corpus. #dialog: total number of dialogs; #turn: average number of turns; #conv.len and #sum.len average dialog and summary number of words respectively. × No summary available.

B Prompt used to generate summary

To generate summary from transcription using ChatGPT-3.5, we used the prompt (translated from French) and one-shot example in table 5.

C Call type classification systems

We report the call type classifier used to generate the call type of the Pipeline_C system and to perform the D_{KL} -based selection in table 6

D All WhisperX size results

We report in the table 7 all the results obtained by training text summarization system on different whisperX size.

Prompt translated from French	
Prompt	<i>Summarize concisely in correct French, without giving too many details, the following text which represents a dialogue between #APPELANT who uses Parisian transport and #CONSEILLER of the RATP. Each turn begins with the name of the speaker in square brackets, followed by the transcription of the turn. Each turn ends with the <END> symbol. Here's the transcript to summarize:</i>
One-shot example (translated from French)	
Dialog	[null] will answer you <END> [agent] good morning <END> [customer] yes good morning madam <END> [customer] madam I'd like to know if tomorrow morning there will be any disruption on the B line of the RER towards Roissy - Charles-de-Gaulle <END> [agent] so the disruptions on the RER B are only in the evening at the airport but not during the day <END> [customer] ah ah are there any disruptions planned for tomorrow evening ? <END> [agent] uh yes because you have works on the section <END> [agent] SNCF so <END> [customer] uh well without that <END> [customer] it works normally? <END> [agent] so if you want that's why but if you take you you leave at what time? <END> [agent] in the morning? <END> [customer] oh <END> [customer] very early <END> [agent] then in that case no <END> [customer] uh ve around 6 o'clock in the <END> [customer] morning <END> [agent] no no there's no problem <END> [customer] there's no problem <END> [agent] no no no no <END> [customer] it's on E only that there are disruptions? <END> [agent] then on the E yes you have d yes yes quite where there it's uh it's also under construction <END> [agent] there's a lot of <END> [customer] <END> [agent] work currently yes <END> [customer] good! in other words uh p tomorrow oh uh no problem on the B <END> [agent] hm no <END> [customer] thank you very much <END> [agent] but I'll <END> [agent] please goodbye sir <END> [customer] goodbye madam <END>.
Summary	A caller wishes to know whether RER Line B will be disrupted the following day, due to work on the route. The advisor says that disruption is expected the next day, but only in the evening.

Table 5: Prompt and One-shot example for ChatGPT 3.5

System	Acc.	W-F1
Conv-CT-classifier	81	80
Syn-CT-classifier	80	80

Table 6: Accuracy and Weighted F1 Score of call types classifiers

Transcription	WER	ROUGE-L	BERTScore	CT-Acc	NE-P	NE-R
Manual	0.0	29.1	38.9	0.77	0.54	0.48
WhisperX Tiny	76.4	24.5	34.4	0.75	0.35	0.28
WhisperX Base	58.6	26.2	35.9	0.75	0.39	0.30
WhisperX Small	46.4	27.0	36.9	0.75	0.49	0.35
WhisperX Medium	40.9	27.6	36.8	0.79	0.51	0.35
WhisperX Large	40.0	27.7	37.1	0.76	0.53	0.36

Table 7: WER score on automatic transcription from different size of WhisperX. ROUGE-L and BERT-scores on summaries generated by Barthez finetuned and evaluated on automatic transcriptions. CT-Acc, NE-P, NE-R, and NE-F1 are our semantic measures based on call type and named entities.

E Sampling parameters

- Top-P: $P \in [0.70, 0.95[$ with the step of 0.05
- Top-K: $K \in [30, 100[$ with the step of 15
- Temperature: $T \in [0.7, 1]$ with the step of 0.1
- Greedy decoding
- Beam decoding: We used a beam of size 6 and keep the 6-best.

F Selection results

We report in table 8 all the scores obtained for summary selection. We include here, the result obtained when using the whisperX base size as automatic transcription system.

System	Criteria	ROUGE-L	BERTScore	CT-Acc	NE-P	NE-R	NE-F1
BARThez + Pipeline _C	BEAM = 4	27.6	37.0	0.79	0.54	0.36	0.43
	min NEHR	27.7	36.4	0.81	0.58	0.38	0.46
	min D_{KL}	27.8	36.8	0.82	0.52	0.38	0.44
	min NEHR + D_{KL}	27.9	37.0	0.82	0.51	0.39	0.44
BARThez + WhisperX Base	BEAM = 4	26.2	35.9	0.75	0.39	0.30	0.34
	min NEHR	25.7	35.5	0.75	0.45	0.40	0.42
	min D_{KL}	24.8	33.0	0.79	0.40	0.31	0.35
	min NEHR + D_{KL}	25.5	35.5	0.77	0.46	0.43	0.44

Table 8: Evaluation of summary selection method. For the model BARThez + Pipeline_C the dataset: DECODA-v2-auto with Whisper Large. For BARThez + WhisperX Base, the dataset DECODA-v2-auto with transcription from WhisperX Base

References

- Ziqiang Cao, Furu Wei, Wenjie Li, and Sujian Li. Faithful to the original: Fact aware neural abstractive summarization. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1), Apr. 2018. doi:10.1609/aaai.v32i1.11912. URL <https://ojs.aaai.org/index.php/AAAI/article/view/11912>.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. On faithfulness and factuality in abstractive summarization. In Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault, editors, *ACL:2020:main*, pages 1906–1919, Online, July 2020. acl. doi:10.18653/v1/2020.acl-main.173. URL [anth#2020.acl-main.173](#)}.
- Tobias Falke, Leonardo F. R. Ribeiro, Prasetya Ajie Utama, Ido Dagan, and Iryna Gurevych. Ranking generated summaries by correctness: An interesting but challenging application for natural language inference. In Anna Korhonen, David Traum, and Lluís Màrquez, editors, *ACL:2019:1*, pages 2214–2220, Florence, Italy, July 2019. acl. doi:10.18653/v1/P19-1213. URL [anth#P19-1213](#)}.
- Feng Nan, Ramesh Nallapati, Zhiguo Wang, Cicero Nogueira dos Santos, Henghui Zhu, Dejiao Zhang, Kathleen McKeown, and Bing Xiang. Entity-level factual consistency of abstractive text summarization. In Paola Merlo, Jorg Tiedemann, and Reut Tsarfaty, editors, *EACL:2021:main*, pages 2727–2733, Online, April 2021. acl. doi:10.18653/v1/2021.eacl-main.235. URL [anth#2021.eacl-main.235](#)}.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12):1–38, March 2023. ISSN 1557-7341. doi:10.1145/3571730. URL <http://dx.doi.org/10.1145/3571730>.
- Esin Durmus, He He, and Mona Diab. FEQA: A question answering evaluation framework for faithfulness assessment in abstractive summarization. In Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault, editors, *ACL:2020:main*, pages 5055–5070, Online, July 2020. acl. doi:10.18653/v1/2020.acl-main.454. URL [anth#2020.acl-main.454](#)}.
- Chih-Wen Goo and Yun-Nung Chen. Abstractive dialogue summarization with sentence-gated modeling optimized by dialogue acts. In *2018 IEEE Spoken Language Technology Workshop (SLT)*, pages 735–742, 2018. doi:10.1109/SLT.2018.8639531.
- Jia Jin Koay, Alexander Roustai, Xiaojin Dai, Dillon Burns, Alec Kerrigan, and Fei Liu. How domain terminology affects meeting summarization performance. In Donia Scott, Nuria Bel, and Chengqing Zong, editors, *COLING:2020:main*, pages 5689–5695, Barcelona, Spain (Online), December 2020. International Committee on Computational Linguistics. doi:10.18653/v1/2020.coling-main.499. URL [anth#2020.coling-main.499](#)}.
- Bin Wang, Chen Zhang, Yan Zhang, Yiming Chen, and Haizhou Li. Analyzing and evaluating faithfulness in dialogue summarization. In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang, editors, *EMNLP:2022:main*, pages 4897–4908, Abu Dhabi, United Arab Emirates, December 2022. acl. doi:10.18653/v1/2022.emnlp-main.325. URL [anth#2022.emnlp-main.325](#)}.

- Xiangru Tang, Arjun Nair, Borui Wang, Bingyao Wang, Jai Desai, Aaron Wade, Haoran Li, Asli Celikyilmaz, Yashar Mehdad, and Dragomir Radev. CONFIT: Toward faithful dialogue summarization with linguistically-informed contrastive fine-tuning. In Marine Carpuat, Marie-Catherine de Marneffe, and Ivan Vladimir Meza Ruiz, editors, *NAACL:2022:main*, pages 5657–5668, Seattle, United States, July 2022. acl. doi:10.18653/v1/2022.naacl-main.415. URL anth#{2022.naacl-main.415}.
- Frederic Bechet, Benjamin Maza, Nicolas Bigouroux, Thierry Bazillon, Marc El-Bèze, Renato De Mori, and Eric Arbillot. DECODA: a call-centre human-human spoken conversation corpus. In Nicoletta Calzolari, Khalid Choukri, Thierry Declerck, Mehmet Uğur Doğan, Bente Maegaard, Joseph Mariani, Asuncion Moreno, Jan Odijk, and Stelios Piperidis, editors, *LREC:2012:1*, pages 1343–1347, Istanbul, Turkey, May 2012. European Language Resources Association (ELRA). URL http://www.lrec-conf.org/proceedings/lrec2012/pdf/684_Paper.pdf.
- Charles T. Hemphill, John J. Godfrey, and George R. Doddington. The ATIS spoken language systems pilot corpus. In *HLT:1990:1*, 1990. URL anth#{H90-1021}.
- Laurence Devillers, Hélène Maynard, Sophie Rosset, Patrick Paroubek, Kevin McTait, D. Mostefa, Khalid Choukri, Laurent Charnay, Caroline Bousquet, Nadine Vigouroux, Frédéric Béchet, Laurent Romary, Jean-Yves Antoine, J. Villaneau, Myriam Vergnes, and J. Goulian. The French MEDIA/EVALDA project: the evaluation of the understanding capability of spoken language dialogue systems. In Maria Teresa Lino, Maria Francisca Xavier, Fátima Ferreira, Rute Costa, and Raquel Silva, editors, *LREC:2004:1*, Lisbon, Portugal, May 2004. European Language Resources Association (ELRA). URL <http://www.lrec-conf.org/proceedings/lrec2004/pdf/356.pdf>.
- S. Kullback and R. A. Leibler. On Information and Sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, 1951. ISSN 0003-4851. URL <https://www.jstor.org/stable/2236703>. Publisher: Institute of Mathematical Statistics.
- Eunice Akani, Benoit Favre, Frederic Bechet, and Romain Gemignani. Reducing named entity hallucination risk to ensure faithful summary generation. In C. Maria Keet, Hung-Yi Lee, and Sina Zarrieß, editors, *INLG:2023:main*, pages 437–442, Prague, Czechia, September 2023. acl. doi:10.18653/v1/2023.inlg-main.33. URL anth#{2023.inlg-main.33}.
- Max Bain, Jaesung Huh, Tengda Han, and Andrew Zisserman. Whisperx: Time-accurate speech transcription of long-form audio. *INTERSPEECH 2023*, 2023.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. Robust speech recognition via large-scale weak supervision, 2022.
- Moussa Kamal Eddine, Antoine Tixier, and Michalis Vazirgiannis. BARThez: a skilled pretrained French sequence-to-sequence model. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih, editors, *EMNLP:2021:main*, pages 9369–9390, Online and Punta Cana, Dominican Republic, November 2021. acl. doi:10.18653/v1/2021.emnlp-main.740. URL anth#{2021.emnlp-main.740}.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need, 2017.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault, editors, *ACL:2020:main*, pages 7871–7880, Online, July 2020. acl. doi:10.18653/v1/2020.acl-main.703. URL anth#{2020.acl-main.703}.
- Louis Martin, Benjamin Muller, Pedro Javier Ortiz Suárez, Yoann Dupont, Laurent Romary, Éric de la Clergerie, Djamé Seddah, and Benoît Sagot. CamemBERT: a tasty French language model. In Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault, editors, *ACL:2020:main*, pages 7203–7219, Online, July 2020. acl. doi:10.18653/v1/2020.acl-main.645. URL anth#{2020.acl-main.645}.