

Effectiveness of French Language Models on Abstractive Dialogue Summarization Task

Yongxin Zhou, François Portet, Fabien Ringeval

Univ. Grenoble Alpes, CNRS, Grenoble INP, LIG

38000 Grenoble, France

{yongxin.zhou, francois.portet, fabien.ringeval}@univ-grenoble-alpes.fr

Abstract

Pre-trained language models have established the state-of-the-art on various natural language processing tasks, including dialogue summarization, which allows the reader to quickly access key information from long conversations in meetings, interviews or phone calls. However, such dialogues are still difficult to handle with current models because the spontaneity of the language involves expressions that are rarely present in the corpora used for pre-training the language models. Moreover, the vast majority of the work accomplished in this field has been focused on English. In this work, we present a study on the summarization of spontaneous oral dialogues in French using several language specific pre-trained models: BART_{thez}, and BelGPT-2, as well as multilingual pre-trained models: mBART, mBART_{thez}, and mT5. Experiments were performed on the DECODA (Call Center) dialogue corpus whose task is to generate abstractive synopses from call center conversations between a caller and one or several agents depending on the situation. Results show that the BART_{thez} models offer the best performance far above the previous state-of-the-art on DECODA. We further discuss the limits of such pre-trained models and the challenges that must be addressed for summarizing spontaneous dialogues.

Keywords: dialogue summarization, pre-trained models, conversation summarization, Non-English summarization

1. Introduction

The task of automatic text summarization consists of presenting textual content in a condensed version that retains the essential information. Recently, the document summarization task has seen a sharp increase in performance due to pre-trained contextualized language models (Liu and Lapata, 2019). Unlike documents, conversations involve multiple speakers, are less structured, and are composed of more informal linguistic usage (Sacks et al., 1978).

Research on dialogue summarization have considered different domains such as summarization of meetings (Buist et al., 2004; Wang and Cardie, 2011; Oya et al., 2014) and phone calls / customer service (Tamura et al., 2011; Stepanov et al., 2015; Liu et al., 2019a).

The introduction of the Transformer architecture (Vaswani et al., 2017) has encouraged the development of large-scale pre-trained language models. Compared to RNN-based models, transformer-based models have deeper structure and more parameters, they use self-attention mechanism to process each token in parallel and represent the relation between words. Such model have pushed the performances forward in automatic summarization. For instance, on the popular benchmark corpus CNN/Daily Mail (Hermann et al., 2015), (Liu and Lapata, 2019) explored fine-tuning BERT (Devlin et al., 2019) to achieve state-of-the-art performance for extractive news summarization, and BART (Lewis et al., 2020) has also improved generation quality on abstractive summarization. However, for dialogue summarization, the impact of pre-trained models is still not sufficiently documented. In particular, it is unclear if the pre-training dataset mismatch

will negatively impact the summarization capability of these models.

In terms of resources, the dialogue summarization task, has been using well-known corpora such as CNN-DailyMail (Hermann et al., 2015), Xsum (Narayan et al., 2018), SAMSum chat-dialogues Corpus (Gliwa et al., 2019) and AMI Meeting Corpus (Carletta et al., 2006). However, most of the available data sets are English language corpora.

In this paper, we present a study to evaluate the impact of pre-trained language models on a dialogue summarization task for the French language. We had evaluated several pre-trained models of French (BART_{thez}, BelGPT-2) and multilingual pre-trained models (mBART, mBART_{thez}, mT5). The experiments were performed on the DECODA (Call Center) dialogue corpus whose task is to generate abstractive synopses from call centre conversations between a caller and one or more agents depending on the situation. Our experimental results establish a new state-of-the-art with the BART_{thez} models on DECODA. The results show that all BART-based pre-trained models have advanced the performance, while the performance of mT5 and BelGPT-2 is not satisfactory enough.

2. Related Work

2.1. Dialogue Summarization

As for summarization tasks, there are two different methods: extractive methods and abstractive methods. Extractive methods consist of copying directly content from the original text, usually sentences. Abstractive methods are not restricted to simply selecting and re-arranging content from the original text, they are able to

generate new words and sentences.

With the development of dialogue systems and natural language generation techniques, the resurgence of dialogue summarization has attracted significant research attention, which aims to condense the original dialogue into a shorter version covering salient information (Feng et al., 2021a). Contrary to well-formed documents, transcripts of multi-person meetings / conversations have various dialogue turns and possible extended topics, summaries generated by NLG models may thus be unfocused on discussion topics.

(Feng et al., 2021a) provide an overview of publicly available research datasets and summarize existing works as well as organize leaderboards under unified metrics, for example, leaderboards of meeting summarization task on AMI (Carletta et al., 2006) and ICSI (Janin et al., 2003), as well as a leaderboard of chat summarization task on SAMSum (Gliwa et al., 2019). Summarized by them, 12 out of 15 major datasets for dialogue summarization are in English.

Taking the chat summarization task on SAMSum dataset as an example, the results showed that the Pre-trained Language Model-based Methods are skilled at transforming the original chat into a simple summary realization (Feng et al., 2021a): the best ROUGE-2 score of the Pre-trained Language Model-based Methods is 28.79 (Feng et al., 2021b), compared to the best ROUGE-2 score of Abstractive Methods at 19.15 (Zhao et al., 2020) and the best ROUGE-2 score of Extractive Methods at 10.27 (LONGEST-3), deep neural models and especially pre-trained language models have significantly improved performance.

According to the domain of input dialogue, summarization on DECODA call-center conversations is of customer service summarization type, and it is task-oriented. Previous work has been done on the DECODA dialogue summarization task: (Trione, 2014; Trione, 2017) used extractive methods and (Stepanov et al., 2015) used abstractive approaches. (Stepanov et al., 2015) used domain knowledge to fill hand-written templates from entities detected in the transcript of the conversation, using topic-dependent rules. Although pre-trained models have shown superior performance on English, however we have few information about how such models would behave for other languages.

2.2. French Language Models and Multilingual Models

If the majority of language models were pre-trained on English-language texts, there have recently been the release of many French oriented pre-trained language models. Furthermore, multilingual pre-trained models have also emerged and showed very good generalisation capability in particular with under-resource languages. A brief summary of the French language models and multilingual models that we used can be found in Table 1.

Regarding French pre-trained models, shortly follow-

ing the release of the BERT-like models: CamemBERT (Martin et al., 2020) and FlauBERT (Le et al., 2020), BARThez (Kamal Eddine et al., 2021) was proposed. It is the French equivalent of the BART model (Lewis et al., 2020) which uses a full transformer architecture (encoder-decoder). Being based on BART, BARThez is particularly well-suited for generative tasks.

For generation, there are also French GPT-like models. For instance, BelGPT-2 (Louis, 2020), which is a GPT-2 model pre-trained on a very large and heterogeneous French corpus (~60Gb). Other French GPT-like models have been proposed in the literature such as PAGnol (Launay et al., 2021) which is a collection of large French language models, GPT-fr (Simoulin and Crabbé, 2021) which uses the OpenAI GPT and GPT-2 architectures (Radford et al., 2018) or even the commercial model Cedille (Müller and Laurent, 2022). We used BelGPT-2 in our experiment since it is open-source and a good trade-off between lightweight model and training size.

Another way to address non-English language tasks is to benefit from multilingual models that were pre-trained on datasets covering a wide range of languages. For example, mBERT (Devlin et al., 2019), mBART (Liu et al., 2020a), XLM-R (Conneau et al., 2020), and mT5 (Xue et al., 2021) are popular models of this type, which are respectively the multilingual variants of BERT (Devlin et al., 2019), BART (Lewis et al., 2020), RoBERTa (Liu et al., 2019b) and T5 (Raffel et al., 2020).

In this study, we will consider mBART, mT5, as well as mBARThez. The latter was designed by the same authors as BARThez (Kamal Eddine et al., 2021) by fine-tuning a multilingual BART on the BARThez training objective, which boosted its performance on both discriminative and generative tasks. mBARThez is thus the most French-oriented multilingual model, this is why we have included it in our experiments.

3. Summarizing Call Center Dialogues in French

3.1. Motivation

As pre-trained models showed advanced performance on summarization tasks in English, we would like to explore the behavior of French pre-trained models on dialogue summarization task. In order to deal with it, we choose the corpus DECODA (Bechet et al., 2012) and the task of Call Centre Conversation Summarization, which had been set during Multiling 2015 (Favre et al., 2015). The objective is to generate abstractive synopses from call center conversations between a caller and one or more agents.

Previous works on DECODA dialogue summarization reported results of extractive methods and abstractive results, our hypothesis is that by using pre-trained language models, we could get better results in this task. Besides, in call center conversations scenarios, the caller (or user) might express their impatience, com-

Model	Description	Pre-training corpus	#params
BARThez(Kamal Eddine et al., 2021)	sequence to sequence pre-trained model	French part of CommonCrawl, NewsCrawl, Wikipedia and other smaller corpora (GIGA, ORTOLANG, MultiUn, EU Bookshop)	165M (architecture: BASE, layers: 12)
BelGPT2(Louis, 2020)	a GPT-2 model pre-trained on a very large and heterogeneous French corpus (~60Gb)	CommonCrawl, NewsCrawl, Wikipedia, Wikisource, Project Gutenberg, EuroParl, NewsCommentary	Small (124M)
mbart.CC25(Liu et al., 2020b)	mBART model with 12 encoder and decoder layers	trained on 25 languages' monolingual corpus, Common Crawl (CC25)	610M
mBARThez(Kamal Eddine et al., 2021)	continue the pretraining of a multilingual BART on BARThez' corpus	the same as BARThez' corpus	458M (architecture: LARGE, layers: 24)
mT5(Xue et al., 2021)	a multilingual variant of T5, Encoder-decoder	new Common Crawl-based dataset covering 101 languages, Common Crawl (mC4)	300M – 13B

Table 1: Summary of French Language Models and Multilingual Models.

plaint and other emotions, which are especially interesting and make DECODA corpus a seldom corpus for further research related to emotional aspects of dialogue summarization.

3.2. DECODA Corpus and Data Partitioning

DECODA call center conversations are a kind of task-oriented dialogue. Conversations happened between the caller (the user) and one agent, sometimes there were more than one agent. The synopses should therefore contain information of users' need and if their problems have been answered or solved by the agent. Top 10 most frequent topics on the DECODA corpus are: Traffic info (22.5%), Directions (17.2%), Lost and found (15.9%), Package subscriptions (11.4%), Timetable (4.6%), Tickets (4.5%), Specialized calls (4.5%), No particular subject (3.6%), New registration (3.4%) and Fare information (3.0%) (Trione, 2014). Besides, a call lasts between 55 seconds for the shortest to 16 minutes for the longest.

Table 2 shows an excerpt of a dialogue. Prefix A (resp B) indicate the employee (resp customer)'s turn.

We used the same data as in the Multiling 2015 CCCS task (Call Centre Conversation Summarization) (Favre et al., 2015): they provided 1000 dialogues without synopses (auto_no_synopses) and 100 dialogues with corresponding synopses for training. As for test data, there are in total 100 dialogues. We divided randomly the original 100 training data with synopses into training and validation sets, with separately 90 and 10 samples. Note that for each dialogue, the corresponding synopsis file each may contain up to 5 synopses writ-

ten by different annotators. Data statistics are presented in Table 3.

3.3. Processing Pipeline

The overall process for dialogue summarization using pre-trained language models is depicted in Figure 1. Pre-trained languages models have limits to the maximum input length. For instance, BART and GPT-2 based models only accept a maximum input length of 1024 tokens, which pose problems when input sequences are longer than their limits. How to process long dialogue summarization is an ongoing research topic (Zhang et al., 2021). In this work, since the objective is to compare pre-trained language models, we use the simple technique of truncation in order to meet the length requirement of pre-trained models. Although truncation is not the most advanced approach, in summarization, using the first part of the document is known to be a strong method and hard to beat, in particular in news where the main information is conveyed by the first line. In the case of call center conversations, we can also assume that the most important information is present at the beginning of the conversation since the client initiates the call by stating his / her problems.

3.3.1. Pre-processing

Since DECODA was acquired in a real and noisy environment, the DECODA transcripts contain various labels related to the speech context, such as <noise.b/>, <noise.rire/>, <noise.it/> *speech* <noise.it/>, etc. There are also <overlap> *speech* </overlap>, which indicates the speech overlap of two speakers. Since the

synopses to make full use of available data for training. In this way, we have 208 training examples, 23 validation examples and 212 test examples.

In experiments, we fine-tuned separately pre-trained models and then used models fine-tuned on DECODA summarization task to generate synopses for test examples. When input sequences exceed the maximum length of the model, we performed truncation as well.

4. Experiments

4.1. Experimental Setup

For all pre-trained models, we used their available models in Huggingface. Experiments for each model were run separately on two versions of the data (*Spk* and *NoSpk*), running on 1 NVIDIA Quadro RTX 8000 48Go GPU server.

For BARThez, mBART, mBARThez and mT5, we used the default parameters for summarization tasks, which are: initial learning rate of $5e-5$, with a `train_batch_size` and `eval_batch_size` of 4, seed of 42. We used Adam optimizer with a linear learning scheduler.

In details, we fine-tuned BARThez¹ (base architecture, 6 encoder and 6 decoder layers) for 10 epochs and saved the best one having the lowest loss on the dev set. Each training took approximately 25 minutes.

mBART² (large architecture, 12 encoder and decoder layers) was also fine-tuned for 10 epochs, each training took approximately 46 minutes.

Regarding mBARThez³, it is also a large architecture with 12 layers in both encoder and decoder. We fine-tuned it for 10 epochs, each training took approximately 38 minutes.

As for mT5⁴ (mT5-Small, 300M parameters), we used `source_prefix: "summarize: "` as a prompt to make the model generate synopses. mT5 was fine-tuned for 40 epochs, each training took approximately 65 minutes. Besides, BelGPT-2 (Louis, 2020) was fine-tuned for 14 epochs, with a learning rate of $5e-5$, `gradient_accumulation_steps` of 32. To introduce summarization behavior, we add `<sep.token>` after the dialogue and set the maximum generated synopsis length at 80 tokens. Each training took approximately 4 hours. During synopsis generation, we tried nucleus sampling and beam search, we reported results with `top_k = 10`, `top_p = 0.5` and `temperature = 1`, which were slightly better.

4.2. Evaluation

During CCCS 2015 (Favre et al., 2015), different systems were evaluated using the CCCS evaluation kit, with ROUGE-1.5.5 and python2.

In our work, models were evaluated with the standard metric ROUGE score (with stemming) (Lin, 2004).

We reported ROUGE-1, ROUGE-2 and ROUGE-L⁵, which calculate respectively the word-overlap, bigram-overlap and longest common sequence between generated summaries and references.

4.3. Quantitative Results

Results on the different pre-trained language models are shown in Table 4.

The table reports the ROUGE-2 scores (Recall only) of the baseline and submitted systems on French DECODA for the multiling challenge (Favre et al., 2015). The first baseline was based on Maximal Marginal Relevance (Baseline-MMR), the second baseline was the first words of the longest turn in the conversation, up to the length limit (Baseline-L), while the third baseline is the words of the longest turn in the first 25% of the conversation, which usually corresponds to the description of the caller’s problem (Baseline-LB). Those baselines are described in more details in (Trione, 2014). Regarding previous computational models, the LIA-RAG system (Linhares et al., 2018) was based on Jensen-Shannon (JS) divergence and TF-IDF approach. We did not find any information about the NTNU system. We did not find more recent results or study on the DECODA dataset.

Overall, in the challenge, the baselines were very difficult to beat with a ROUGE-2 of 4.0. This quite low score shows the great difficulty of the task. Indeed, the transcriptions contain overlapping turns, fillers, pauses, noise, partial words, etc. Furthermore translators were required to translate speech phenomena such as disfluencies as closely as possible to the source language maintaining ‘naturalness’ in the target language.

Comparing the performance of different pre-trained models fine-tuned on the DECODA summarization task, BARThez outperformed other models, with a ROUGE-2 score of 16.85 using the *Spk* data version.

In general, all pre-trained models outperform previous results. BART-based pre-trained models (BARThez, mBART, mBARThez) have better performance than mT5 and BelGPT-2.

The difference between *Spk* and *NoSpk* is not significant. For some models like BARThez and mBART, *Spk* (using speakers’ prefix) performs slightly better than *NoSpk*; while for mT5, *NoSpk* (without the speakers’ prefix) performs better than *Spk*.

4.4. Qualitative Analysis

We observed that the BARThez pre-trained model fine-tuned on the *Spk* version of DECODA led to the best ROUGE. We then randomly selected 20 test samples and performed an error analysis of the generated predictions, comparing them to references.

¹<https://huggingface.co/moussaKam/barthez>

²<https://huggingface.co/facebook/mbart-large-cc25>

³<https://huggingface.co/moussaKam/mbarthez>

⁴<https://huggingface.co/google/mt5-small>

⁵The ROUGE metric used in Hugging Face is a wrapper around Google Research reimplementation of ROUGE: <https://github.com/google-research/google-research/tree/master/rouge>, we used it to report all our results.

#	Models	ROUGE-1	ROUGE-2	ROUGE-L
(Favre et al., 2015)	Baseline-MMR	-	4.5	-
(Favre et al., 2015)	Baseline-L	-	4.0	-
(Favre et al., 2015)	Baseline-LB	-	4.6	-
(Favre et al., 2015)	NTNU:1	-	3.5	-
(Linhares et al., 2018)	LIA-RAG:1	-	3.7	-
(Kamal Eddine et al., 2021)	BARThez #Spk	35.45	16.85	29.45
(Kamal Eddine et al., 2021)	BARThez #NoSpk	33.01	15.30	27.92
(Louis, 2020)	BelGPT-2 #Spk	17.73	5.15	13.82
(Louis, 2020)	BelGPT-2 #NoSpk	18.06	4.55	13.15
(Liu et al., 2020a)	mBART #Spk	33.76	14.13	27.66
(Liu et al., 2020a)	mBART #NoSpk	31.60	13.69	26.20
(Kamal Eddine et al., 2021)	mBARThez #Spk	34.95	15.86	28.85
(Kamal Eddine et al., 2021)	mBARThez #NoSpk	35.31	15.52	28.61
(Xue et al., 2021)	mT5 #Spk	23.90	8.69	20.06
(Xue et al., 2021)	mT5 #NoSpk	27.82	11.08	23.36

Table 4: ROUGE-2 performances of the baseline and submitted systems for the multiling challenge 2015 on French DECODA and results of ROUGE-1, ROUGE-2 and ROUGE-L scores of the pre-trained language models. The best results are boldfaced.

FR_20091112_RATP_SCD_0826	
References	Request for train schedules from Meaux station to Gare de l’Est at a given time. (1) schedules RER E from Meaux to Gare de l’Est (166)
Spk	
BARThez	Request for train schedules to go from Meaux station to Paris.
BelGPT-2	Request for information on the destination of the Meaux-Chappes-Saint-Lazare subway line. Bus transfer at Saint-Lazare station. Is the call possible. Is the call possible to do it. Is the call possible. Is the call possible. Is the call
mBART	Request for SNCF train schedules to go to Paris from the station of Meaux from the station of Gare de Lyon in Paris. Request for timetables to go to the station of Meaux at sixteen hours forty-nine by leaving the station of Meaux at sixteen hours forty-nine by leaving the station of Gare de Lyon.
mBARThez	train schedules in Île-de-France
mT5	Request for timetable to go to the station of Meaux.
NoSpk	
BARThez	Request for train schedules to go to Paris.
BelGPT-2	, the journey from the station of Meaux to the Gare de Lyon, passing by the station of Bercy) and finally the journey from the station of Bercy to the station of Bercy, passing by the station of Bercy, the journey from the station of Bercy to the station of Bercy, and finally the journey from the station of Bercy to the station of Montparnasse, the journey from the station of Bercy
mBART	Request for schedules for an RER line between Meaux Meaux station and Gare de Lyon station. Contact with the concerned service.
mBARThez	RER B schedules to go to Paris from Meaux station
mT5	Request for timetable to go to the station of Meaux.

Table 5: (Translation in English) Doc *FR_20091112_RATP_SCD_0826* from DECODA’s test set, and associated reference and generated summaries by fine-tuning different models. Errors are marked in red.

The references were written by different annotators and not all of them summarized the dialogue in the same way. Furthermore, because synopses are used to quickly convey information about the conversation between the caller and the RATP agent, they are even

more abstract than typical dialogue summaries. After comparing the synopses generated by BARThez (*Spk*) to the references written by annotators, we observed several common types of errors:

- **Hallucination:** This is a common type of error in

neural NLG and is widely present in our generated synopses, meaning that the generated synopses are inconsistent or unrelated to the system input.

- **Omission:** The other most common type of error, which means that some of the information presented in the references is not presented in the generated synopses.
- **Grammatical Error:** In our case, the generated synopses are generally smooth, but they may have incorrect syntax or semantics.

We categorize the errors and count their frequencies in Table 6.

Type of errors	Occurrences
Hallucination	15 / 20
Omission	15 / 20
Grammatical error	5 / 20

Table 6: The most common error types of the BARThez (*Spk*) model compared to the golden reference over 20 sampled dialogues, with the number of occurrences for each error type.

We noticed that 7 of the 20 generated synopses have a *Communication of the relevant service number* at the end, which is not presented in the reference or in the input dialogue. Since we truncated the length of the input sequences to a maximum of 1024 tokens, this could be the result of truncation: the lack of information in the final part of the dialogue causes system to sometimes generate a general sentence.

In addition, for each synopsis generated, we counted the number of errors. Note that for the same type of errors, we can count more than once. For example, for omission, if two different points of information are missing from the generated synopsis, we count twice. The results of the number of errors are reported in Table 7, the total number of errors counts from 1 to 5 for above 20 randomly selected synopses.

Number of errors	Synopses (total: 20)
1	4
2	3
3	10
4	2
5	1

Table 7: Number of synopses generated by BARThez (*Spk*) with a total number of errors from 1 to 5.

We noticed that half of the synopses generated contain 3 errors, 1 sample contains 5 errors and 4 synopses contain only 1 error. For example, compared with the reference *demande de précisions sur la tarification du métro, mais l'appelant, un particulier, est passé directement par le numéro VGC, refus du conseiller de*

le prendre en charge comme grand compte, indication du numéro 3246 pour qu'il rappelle, the generated synopsis *Demande de renseignement sur les tarifs RATP sur Paris ou la zone parisienne. Communication du numéro de téléphone.* has 3 errors :

- 1) omission - *mais l'appelant, un particulier, est passé directement par le numéro VGC;*
- 2) omission - *refus du conseiller de le prendre en charge comme grand compte;*
- 3) omission - *indication du numéro 3246 pour qu'il rappelle*

Besides, while doing error analysis, we also noticed different writing styles of different annotators. For example, some golden references always begin with *ask for (Demande de)* while some other references are much shorter, for example *lost phone > but not found (perdu téléphone > mais pas retrouvé).*

Taking *FR_20091112_RATP_SCD_0826* from DECODA's test set as an example, we compared the summaries generated by fine-tuning different pre-trained models with the references in Table 5 (translated in English, original French version and the dialogue in Appendix). Errors are marked in red.

For BARThez (*Spk*) which got the highest ROUGE-2 score, the synopsis generated for this test example is quite good, however BARThez (*NoSpk*) does not mention *from Meaux station*.

In general, for fine-tuned BelGPT-2, the generated synopses are correct in some respects but suffer from factual inconsistency. The model started to generate repetitive word sequences quickly. Its poor performance may be due to the type of model (decoder only) and its small size.

mBART (*Spk*) was good in the beginning, but it started to produce incorrect and unrelated information later on. As for mBART (*NoSpk*), it incorrectly predicted *Gare de Lyon station* instead of *Gare de l'EST*.

As for fine-tuned mBARThez (*Spk*), it seems to be correct but missed some information like *from Meaux to Gare de l'Est*. mBARThez (*NoSpk*) wrongly predicted *RER B* instead of *RER E*.

mT5 wrongly predicted the direction, it should be *from Meaux station to Gare de l'Est* but it wrongly took *to go to the station of Meaux*. In addition, the generated synopses lack other information such as *request for train schedules, at a given time* and *to Gare de l'Est*.

5. Discussion and Conclusion

Our experimental results show that the BARThez models provide the best performance, far beyond the previous state of the art on DECODA. However, even though the pre-trained models perform better than traditional extractive or abstractive methods, their outputs are difficult to control and suffer from hallucination and omission, there are also some grammatical errors.

On the one hand, the limit of maximum token length for pre-trained models evokes challenges to summarize such non-trivial dialogue. For long dialogue summarization, we can use retrieve-then-summarize pipeline or use end-to-end summarization models (Zhang et al., 2021). (Zhu et al., 2020) designed a hierarchical structure named HMNet to accommodate long meeting transcripts and a role vector to depict the difference between speakers. (Beltagy et al., 2020) introduced the Longformer with an attention mechanism that scales linearly with sequence length, making it easy to process documents of thousands of tokens or longer. However, these end-to-end summarization models capable of dealing with long sequences problems are both pre-trained in English language.

On the other hand, pre-trained models were mostly learned on news corpora but far from the spontaneous conversation (repetitions, hesitation, etc.). Dialogue summarization therefore remains a challenging task. Besides, in call center conversations scenarios, the caller (or user) might express their impatience, complaint or other emotions, we would like to further explore DEOCDA corpus for emotion-related dialogue summarization.

6. Acknowledgements

This research has received funding from the Banque Publique d'Investissement (BPI) under grant agreement THERADIA, the Association Nationale de la Recherche et de la Technologie (ANRT), under grant agreement No.2019/0729, and has been partially supported by MIAI@Grenoble-Alpes (ANR-19-P3IA-0003).

Appendix

Table 8 presents the references and summaries generated by fine-tuning different models of the test example *FR_20091112_RATP_SCD_0826*, *Spk* and *NoSpk* versions of the dialogue are also shown below.

Spk version: This version presents the conversations with the speakers' information.

<Spk_A> bonjour
 <Spk_B> oui bonjour madame
 <Spk_B> je vous appelle pour avoir des horaires de train en+fait c' est pas pour le métro je sais pas si vous pouvez me les donner ou pas
 <Spk_A> trains SNCF
 <Spk_B> oui trains SNCF oui
 <Spk_A> vous prenez quelle ligne monsieur
 <Spk_B> euh la ligne euh enfin en+fait c' est pas SNCF enfin c' est Île-de-France quoi je sais pas comment ils appellent ça
 <Spk_B> RER voilà c' est pour les RER
 <Spk_B> voilà et euh je prends la ligne euh Meaux de Meaux pour aller à Paris je sais plus c' est laquelle c' est la
 <Spk_B> E je crois
 <Spk_A> d'accord et vous voulez donc les horaires

euh pour quel jour et à quel moment monsieur
 <Spk_B> euh là pour euh tout à l' heure euh pour euh aux environs de dix-sept heures en partant de la gare de Meaux

<Spk_B> euh vers la Gare+de+l'Est à Paris
 <Spk_A> alors vous partez de Meaux et vous allez donc à la Gare+de+l'Est <Spk_A> et vous voudriez les horaires <Spk_B> voilà
 <Spk_A> vers euh dix-sept heures <Spk_A> alors moi je peux regarder ce+que <Spk_B> ouais dix-sept heures ouais
 <Spk_A> j' ai comme horaires un instant monsieur
 <Spk_A> s'il+vous+plaît <Spk_B> d'accord il y a
 <Spk_B> pas de souci
 <Spk_A> monsieur s'il+vous+plaît
 <Spk_B> oui
 <Spk_A> donc voilà ce+que j' ai comme horaires moi vous avez donc un départ à la gare de Meaux donc à seize heures quarante-neuf
 <Spk_B> seize heures quarante-neuf d'accord
 <Spk_A> et après vous avez celui de dix-sept heures dix-neuf
 <Spk_B> alors seize heures quarante-neuf dix-sept heures dix-neuf <Spk_A> OK d'accord <Spk_B> oui
 <Spk_B> ben je vous remercie
 <Spk_A> mais je vous en prie <Spk_B> bonne journée madame
 <Spk_A> merci à vous aussi monsieur au+revoir
 <Spk_B> au+revoir

NoSpk version: This version presents the conversations without the speakers' information.

bonjour. oui bonjour madame. je vous appelle pour avoir des horaires de train en+fait c' est pas pour le métro je sais pas si vous pouvez me les donner ou pas. trains SNCF. oui trains SNCF oui. vous prenez quelle ligne monsieur. euh la ligne euh enfin en+fait c' est pas SNCF enfin c' est Île-de-France quoi je sais pas comment ils appellent ça. RER voilà c' est pour les RER. voilà et euh je prends la ligne euh Meaux de Meaux pour aller à Paris je sais plus c' est laquelle c' est la. E je crois. d'accord et vous voulez donc les horaires euh pour quel jour et à quel moment monsieur. euh là pour euh tout à l' heure euh pour euh aux environs de dix-sept heures en partant de la gare de Meaux. euh vers la Gare+de+l'Est à Paris. alors vous partez de Meaux et vous allez donc à la Gare+de+l'Est et vous voudriez les horaires voilà. vers euh dix-sept heures alors moi je peux regarder ce+que ouais dix-sept heures ouais. j' ai comme horaires un instant monsieur s'il+vous+plaît d'accord il y a. pas de souci. monsieur s'il+vous+plaît. oui. donc voilà ce+que j' ai comme horaires moi vous avez donc un départ à la gare de Meaux donc à seize heures quarante-neuf. seize heures quarante-neuf d'accord. et après vous avez celui de dix-sept heures dix-neuf. alors seize heures quarante-neuf dix-sept heures dix-neuf OK d'accord oui. ben je vous remercie. mais je vous en prie bonne journée madame. merci à vous aussi monsieur au+revoir. au+revoir.

FR_20091112_RATP_SCD_0826	
References	Demande d’horaires de train de la gare de Meaux à la gare de l’Est à une heure donnée. (1) horaires RER E de Meaux à la Gare de l’Est (166)
Spk	
BARThez	Demande d’horaires de train pour aller de la gare de Meaux à Paris.
BelGPT-2	demande de renseignement sur la destination de la ligne de méto de Meaux-Chappes-Saint-Lazare. Transfert de bus à la gare de Saint-Lazare. L’appel est-il possible. L’appel est-il possible de le faire. L’appel est-il possible. L’appel est-il possible. L’appel est
mBART	Demande d’horaires de trains SNCF pour aller à Paris en partant de la gare de Meaux en partant de la gare de Gare de Lyon à Paris. Demande des horaires pour se rendre à la gare de Meaux à seize heures quarante-neuf en partant de la gare de Meaux à seize heures quarante-neuf en partant de la gare de Gare de Lyon.
mBARThez	horaires de train en Île-de-France
mT5	Demande de horaires pour aller à la gare de Meaux.
NoSpk	
BARThez	Demande d’horaires de train pour aller à Paris.
BelGPT-2	, le trajet de la gare de Meaux vers la Gare de Lyon, en passant par la gare de Bercy) et enfin le trajet de la gare de Bercy à la gare de Bercy, en passant par la gare de Bercy, le trajet de la gare de Bercy vers la gare de Bercy, et enfin le trajet de la gare de Bercy vers la gare de Montparnasse, le trajet de la gare de Bercy
mBART	Demande d’horaires pour une ligne de RER entre la gare de Meaux de Meaux et la gare de Gare de Lyon. Mise en relation avec le service concerné.
mBARThez	horaires du RER B pour aller à Paris en partant de la gare de Meaux
mT5	Demande de horaires pour aller à la gare de Meaux.

Table 8: Doc *FR_20091112_RATP_SCD_0826* from DECODA’s test set, and associated reference and generated summaries by fine-tuning different models. Errors are marked in red.

7. Bibliographical References

- Beltagy, I., Peters, M. E., and Cohan, A. (2020). Long-former: The long-document transformer. *CoRR*, abs/2004.05150.
- Buist, A. H., Kraaij, W., and Raaijmakers, S. (2004). Automatic summarization of meeting data: A feasibility study. In *in Proc. Meeting of Computational Linguistics in the Netherlands (CLIN)*.
- Conneau, A., Khandelwal, K., Goyal, N., Chaudhary, V., Wenzek, G., Guzmán, F., Grave, E., Ott, M., Zettlemoyer, L., and Stoyanov, V. (2020). Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online, July. Association for Computational Linguistics.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (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, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Favre, B., Stepanov, E., Trione, J., Béchet, F., and Riccardi, G. (2015). Call centre conversation summarization: A pilot task at multiling 2015. In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 232–236, Prague, Czech Republic, September. Association for Computational Linguistics.
- Feng, X., Feng, X., and Qin, B. (2021a). A survey on dialogue summarization: Recent advances and new frontiers. *CoRR*, abs/2107.03175.
- Feng, X., Feng, X., Qin, L., Qin, B., and Liu, T. (2021b). Language model as an annotator: Exploring DialoGPT for dialogue summarization. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1479–1491, Online, August. Association for Computational Linguistics.
- Kamal Eddine, M., Tixier, A., and Vazirgiannis, M. (2021). BARThez: a skilled pretrained French sequence-to-sequence model. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9369–9390, Online and Punta Cana, Dominican Republic, November. Association for Computational Linguistics.
- Launay, J., Tommasone, G. L., Pannier, B., Boniface, F., Chatelain, A., Cappelli, A., Poli, I., and Seddah, D. (2021). PAGnol: An Extra-Large French Generative Model. Research report, LightON, October.
- Le, H., Vial, L., Frej, J., Segonne, V., Coavoux,

- M., Lecouteux, B., Allauzen, A., Crabbé, B., Becsac, L., and Schwab, D. (2020). FlauBERT: Un-supervised language model pre-training for French. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 2479–2490, Marseille, France, May. European Language Resources Association.
- Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., and Zettlemoyer, L. (2020). BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online, July. Association for Computational Linguistics.
- Lin, C.-Y. (2004). ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain, July. Association for Computational Linguistics.
- Linhares, E., Torres-Moreno, J.-M., and Linhares, A. C. (2018). LIA-RAG: a system based on graphs and divergence of probabilities applied to Speech-To-Text Summarization. working paper or preprint, April.
- Liu, Y. and Lapata, M. (2019). Text summarization with pretrained encoders. 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 3730–3740, Hong Kong, China, November. Association for Computational Linguistics.
- Liu, C., Wang, P., Xu, J., Li, Z., and Ye, J. (2019a). Automatic dialogue summary generation for customer service. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '19*, page 1957–1965, New York, NY, USA. Association for Computing Machinery.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., and Stoyanov, V. (2019b). Roberta: A robustly optimized bert pre-training approach.
- Liu, Y., Gu, J., Goyal, N., Li, X., Edunov, S., Ghazvininejad, M., Lewis, M., and Zettlemoyer, L. (2020a). Multilingual denoising pre-training for neural machine translation. *Transactions of the Association for Computational Linguistics*, 8:726–742.
- Liu, Y., Gu, J., Goyal, N., Li, X., Edunov, S., Ghazvininejad, M., Lewis, M., and Zettlemoyer, L. (2020b). Multilingual denoising pre-training for neural machine translation. arXiv, <https://arxiv.org/abs/2001.08210>.
- Louis, A. (2020). BelGPT-2: a GPT-2 model pre-trained on French corpora. <https://github.com/antoiloui/belgpt2>.
- Martin, L., Muller, B., Ortiz Suárez, P. J., Dupont, Y., Romary, L., de la Clergerie, É., Seddah, D., and Sagot, B. (2020). CamemBERT: a tasty French language model. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7203–7219, Online, July. Association for Computational Linguistics.
- Müller, M. and Laurent, F. (2022). Cedille: A large autoregressive french language model. arXiv, <https://arxiv.org/abs/2202.03371>.
- Oya, T., Mehdad, Y., Carenini, G., and Ng, R. (2014). A template-based abstractive meeting summarization: Leveraging summary and source text relationships. In *Proceedings of the 8th International Natural Language Generation Conference (INLG)*, pages 45–53, Philadelphia, Pennsylvania, U.S.A., June. Association for Computational Linguistics.
- Radford, A., Narasimhan, K., Salimans, T., and Sutskever, I. (2018). Improving language understanding by generative pre-training.
- Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., and Liu, P. J. (2020). Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Sacks, H., Schegloff, E. A., and Jefferson, G. (1978). A simplest systematics for the organization of turn taking for conversation. In *Studies in the organization of conversational interaction*, pages 7–55. Elsevier.
- Simoulin, A. and Crabbé, B. (2021). Un modèle transformer génératif pré-entraîné pour le français. In *Traitement Automatique des Langues Naturelles*, pages 245–254. ATALA.
- Stepanov, E., Favre, B., Alam, F., Chowdhury, S., Singla, K., Trione, J., Béchet, F., and Riccardi, G. (2015). Automatic summarization of call-center conversations. In *Conference: IEEE Automatic Speech Recognition and Understanding Workshop (ASRU 2015) At: Scottsdale, Arizona, USA*, 12.
- Tamura, A., Ishikawa, K., Saikou, M., and Tsuchida, M. (2011). Extractive summarization method for contact center dialogues based on call logs. In *Proceedings of 5th International Joint Conference on Natural Language Processing*, pages 500–508, Chiang Mai, Thailand, November. Asian Federation of Natural Language Processing.
- Trione, J. (2014). Extraction methods for automatic summarization of spoken conversations from call centers (méthodes par extraction pour le résumé automatique de conversations parlées provenant de centres d’appels) [in French]. In *Proceedings of TALN 2014 (Volume 4: RECITAL - Student Research Workshop)*, pages 104–111, Marseille, France, July. Association pour le Traitement Automatique des Langues.
- Trione, J. (2017). *Extraction methods for automatic summarization of spoken conversations from call centers (Méthodes par extraction pour le résumé automatique de conversations parlées provenant de*

centres d'appels) [in French]. Thèse de doctorat en informatique, Aix-Marseille université.

- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., and Polosukhin, I. (2017). Attention is all you need. In I. Guyon, et al., editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Wang, L. and Cardie, C. (2011). Summarizing decisions in spoken meetings. In *Proceedings of the Workshop on Automatic Summarization for Different Genres, Media, and Languages*, pages 16–24, Portland, Oregon, June. Association for Computational Linguistics.
- Xue, L., Constant, N., Roberts, A., Kale, M., Al-Rfou, R., Siddhant, A., Barua, A., and Raffel, C. (2021). mT5: A massively multilingual pre-trained text-to-text transformer. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online, June. Association for Computational Linguistics.
- Zhang, Y., Ni, A., Yu, T., Zhang, R., Zhu, C., Deb, B., Celikyilmaz, A., Awadallah, A. H., and Radev, D. (2021). An exploratory study on long dialogue summarization: What works and what's next. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4426–4433, Punta Cana, Dominican Republic, November. Association for Computational Linguistics.
- Zhao, L., Xu, W., and Guo, J. (2020). Improving abstractive dialogue summarization with graph structures and topic words. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 437–449, Barcelona, Spain (Online), December. International Committee on Computational Linguistics.
- Zhu, C., Xu, R., Zeng, M., and Huang, X. (2020). A hierarchical network for abstractive meeting summarization with cross-domain pretraining. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 194–203, Online, November. Association for Computational Linguistics.
- pre-announcement. In Steve Renals et al., editors, *Machine Learning for Multimodal Interaction*, pages 28–39, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Gliwa, B., Mochol, I., Biesek, M., and Wawer, A. (2019). SAMSum corpus: A human-annotated dialogue dataset for abstractive summarization. In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pages 70–79, Hong Kong, China, November. Association for Computational Linguistics.
- Hermann, K. M., Kocisky, T., Grefenstette, E., Espeholt, L., Kay, W., Suleyman, M., and Blunsom, P. (2015). Teaching machines to read and comprehend. *Advances in neural information processing systems*, 28:1693–1701.
- Janin, A., Baron, D., Edwards, J., Ellis, D., Gelbart, D., Morgan, N., Peskin, B., Pfau, T., Shriberg, E., Stolcke, A., and Wooters, C. (2003). The icisi meeting corpus. In *2003 IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003. Proceedings. (ICASSP '03).*, volume 1, pages I–I.
- Narayan, S., Cohen, S. B., and Lapata, M. (2018). Don't give me the details, just the summary! Topic-aware convolutional neural networks for extreme summarization. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Brussels, Belgium.

8. Language Resource References

- Bechet, F., Maza, B., Bigouroux, N., Bazillon, T., El-Bèze, M., De Mori, R., and Arbillot, E. (2012). DECODA: a call-centre human-human spoken conversation corpus. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 1343–1347, Istanbul, Turkey, May. European Language Resources Association (ELRA).
- Carletta, J., Ashby, S., Bourban, S., Flynn, M., Guillemot, M., Hain, T., Kadlec, J., Karaiskos, V., Kraaij, W., Kronenthal, M., Lathoud, G., Lincoln, M., Lisowska, A., McCowan, I., Post, W., Reidsma, D., and Wellner, P. (2006). The ami meeting corpus: A