Proceedings of the 12th Conference on Language Resources and Evaluation (LREC 2020), pages 574–580
Marseille, 11–16 May 2020
© European Language Resources Association (ELRA), licensed under CC-BY-NC

A French Medical Conversations Corpus Annotated for a Virtual Patient Dialogue System

Fréjus A. A. Laleye[1], Gaël de Chalendar[1], Antonia Blanié[2], Antoine Brouquet[2], Dan Benhamou[2]
1 CEA, LIST, Laboratoire d'Analyse Sémantique Texte et Image, Gif-sur-Yvette, F-91191
2 Laboratoire de Formation par la Simulation et l'Image en Médecine et en Santé (LabForSIMS),
Faculté de Médecine Paris-Sud, 94275 Le Kremlin Bicêtre
{frejus.laleye,gael.de-chalendar}@cea.fr,
{antonia.blanie,antoine.brouquet,dan.behnamou}@aphp.fr

Abstract

Data-driven approaches for creating virtual patient dialogue systems require the availability of large data specific to the language, domain and clinical cases studied. Based on the lack of dialogue corpora in French for medical education, we propose an annotated corpus of dialogues including medical consultation interactions between doctor and patient. In this work, we detail the building process of the proposed dialogue corpus, describe the annotation guidelines and also present the statistics of its contents. We then conducted a question categorization task to evaluate the benefits of the proposed corpus that is made publicly available.

Keywords: Virtual Patient, Medical Corpus, Question Answering System

1. Introduction

Virtual patients allow medical students to perform clinical simulations to acquire experience through the practice of critical procedures without misbehaving with or endangering real patients. For the students to be able to interact naturally with their virtual patient, the dialogue system must be quite versatile. It must have the capacity to understand a wide variety of questions and to answer with natural sounding sentences. Purely pattern-based systems do not have this capacity. That's why they must be completed or replaced by statistical or, more recently, neural-based ones. To train the latter, large corpora of medical interviews are necessary. Several recent works such as (Jin et al., 2017; Maicher et al., 2019; Spänig et al., 2019) have explored the potential advantages of statistical and neural approaches using a large amounts data to enhance the mechanism for interpreting questions in dialogue systems.

In this paper, we present the corpus of medical conversations in French that we have built in the context of the development of such a virtual patient dialogue system. We describe the methodology we used, the annotations we made and several data about its content. We also describe one of the subsystems that were built using the corpus, namely a question categorization tool. The remaining of the dialog system is described and evaluated in (Laleye et al., 2020). We make our corpus freely available under a Free/Libre Open Source licence[1].

2. Related Work

There are different approaches to achieving a dialogue system for a virtual patient system. Depending on the clinical case and the educational objectives, the dialogue management approaches are divided into handcrafted (Laroche et al., 2010) and probabilistic (Celikyilmaz et al., 2018). Instead of defining rules for the dialog strategy by hand, probabilistic dialogue management systems learn appropriate answers from a large corpus by matching the last utterance with an example in the training dataset and uses the response from the training set (Harms et al., 2019). Probabilistic methods such as statistical or neural approaches require the use of large corpora to be applicable to Virtual Patient Systems. Due to the lack of existing dialogue corpora for medical education, several studies (Maicher et al., 2016; Campillos Llanos et al., 2019) have been limited in the use of these approaches.

Similarly to our work, there are efforts to make available resources built from real data for corpus-based data-driven dialogue systems. For the task of training a medical student in the interaction with patients, the datasets differ significantly from the clinical case and the educational objectives point of views. In (Campillos Llanos et al., 2019), the authors designed a dialogue system for handling a wide variety a dialogue specialties and clinical cases. Due to unavailable of dialogue data for the task and domain, they built their own corpus of dialogue from the collected data[2]. 32 computer science students and researchers interacted with their system for the development and 39 medical students and doctors for the evaluation. The authors in (Porhet et al., 2017) developed a virtual patient able to interact in a multimodal way with doctors announcing an undesirable event in order to facilitate the doctor's training to break bad news. In their work, they proposed a corpus[3] composed of 13 videos transcribed and annotated of patient-doctor interaction with different scenarios. The authors in (Gokcen et al., 2016) proposed a corpus of 104 dialogues between early stage medical students and a virtual standardized patient. They have manually annotated word alignments for 942 sentence pairs in order to determine for each question asked by the medical student which of the set of questions

[1] Temporary location during review: `https://github.com/kleag/labforsims2-corpus`.

[2] `https://pvdial.limsi.fr/data/PG-logs-eval.zip`

[3] `http://www2.lpl-aix.fr/~acorformed/17-ICMI/annexe_1.html`

anticipated by the content author best matches the student's question.

3. Methodology

This section provides an overview of our question answering dialogue corpus. It presents the data sources and annotation methods. The corpus is primarily built for tasks such as question answering system, dialog management system of a virtual patient, medical question categorization, clinical reasoning analysis and medical information extraction. The aim is to propose a dataset of questions generally encountered during a surgical emergency consultation for abdominal pain. The dialogues are defined to allow a system to handle complex questions that require several inference steps as in (Iyyer et al., 2017). This feature is very useful for systems that must refer to the entities in previous questions to respond to new ambiguous questions.

3.1. Data Collection

The first step was to identify sources containing data that respond to the studied clinical case. The first source of data consists of two surgeons, authors of this paper, who took care to define the scenario of the clinical case. They initially built a first set of questions and responses with information extracted from a patient's record. This information is structured in sections that include the types of information needed by doctor for diagnosis. The patient's record is presented in Table 1 and describes common data found in medical consultations.

This first dataset containing information defined by physician trainers is made available to a virtual standardized patient (VSP) built to interact with medical students in order to collect more data. The dialogue manager of our VSP is a hand-authored pattern matching system which dialogues with students via a speech recognizer and a speech synthesis module. The clinical record of the virtual patient is presented in Table 2.

We then deployed the VSP in a medical school to collect more data through interactions with interns who read the clinical record prior to consultation. Our second data source consists of 41 interns who interacted with the virtual patient, each having an end-to-end dialogue in the goal to browse all types of information in the patient record. The collected content were manually rearranged and annotated according to the information in the patient record.

3.2. Annotation process

The collected data were first cleaned and then organized into two sets. In the first set named *Single-turn dataset*, we included the question and response pairs taken separately and which do not integrate a dialogue. The second, named *Context QA dataset*, includes all the end-to-end dialogues that are the interactions between medical students and the virtual patient. The first step in the annotation process was to assign one of the categories defined in Table 3 to the questions and responses by referring to the patient's record. Thus, a tag has been associated with each entry of the corpus according to its category. This task was carried out with the help of the expert surgeon doctors with whom this work was done.

3.2.1. *Single-turn dataset* annotation guidelines

This dataset contains data organized into three sets: the doctor's questions, the patient responses and a set that matches a response to a question. A question is identified by the tag of its category followed by its sequence number in the doctor set. Some examples of annotated questions are presented in Table 4 and the list of tags with categories in Table 3 (column 2). A response is identified by the letter **P** followed by its sequence number in the patient set. Table 5 lists some examples of annotated responses. The cross between questions and responses are shown in Table 6.

3.2.2. *Context QA dataset* annotation guidelines

Context QA dataset is organized in dialogues equal to the number of medical students who participated in the data collection. We tried to follow the same annotation guidelines previously detailed by extending the categories to the types of information to be extracted from the patient record as detailed in Table 1. The tag ids were used for category identification (see column 3, Table 3) We have then assigned to each type of information an integer, as an id, written on three characters (e.g *symptom_fever* corresponds to a question about fever in the symptom category with 001 as id; *lifestyle_addictions* corresponds to a question about the patient addictions in the lifestyle category with 000 as id). These detailed type of information has been categorized by ourselves with regard to the content of the dialogues. This overall strategy was therefore used for labeling the questions and responses in the 41 dialogues of the corpus. A question of a dialogue is then identified by:

1. the letter *d* to mean that it's a doctor's sentence;
2. the number assigned to the dialogue (e.g 01 for the first dialogue)
3. the id of its category tag (e.g 01 for personal data);
4. the id of the information type from the category (e.g 005 for the patient's profession);
5. its order in the dialogue (e.g 2 if it appears second in the turns of speech);

The same identification is applied to the response associated with a question except that its id starts with the letter *p*. Table 7 shows an example of an annotated dialog extracted from the corpus with description of the ids

4. Corpus structure and statistics

The corpus is organized in interactions of a real medical consultation. A dialogue represents a consultation containing patient responses associated with the doctor questions and related to an information context. A question is structured so as to easily find its type and its response. The annotated questions take into account variations due to either the question type either the structure of each question type. Dialogues are natural and user-friendly. Some language phenomena have been taken into account in building the corpus. These phenomena are:

the structural variation of the question illustrated by the use of different formulations of the same question in

Aim of consultation	Personal Data	Medical History	Symptoms	Lifestyle	Treatments
The Goal	First name Last name Age Weight Housing Job Children	Family history Past medical history Past surgical history Allergies Medications taken	Sickness history Changes/evolutions Location Timing/chronologie Quantification/severity	Addictions Pets	Type of treatment Method Date and period Observation

Table 1: Patient record.

```
aimOfConsultation:
   Goal : consultation for abdominal pain
Personal data:
   Fisrt name: Simone
   Last name : Labforsims
   Age       : 42
   Weight    : 82
   Housing   : she lives with family in
                Gif/yvette
    Job      : teacher in kindergarten
    Children : she has 2 children and she
                miscarried in 97
Medical history:
   Family history: her father had a colon
                   cancer at age 70, and
                   he had a heart attack
                   last year
    Past medical history: no particular
                   disease
    Past surgical history: operation of
                   wisdom teeth
    Allergies : allergic pollen
    Madications taken: paracetamol occasio-
                nally
Symptom:
    Sickness history: the pain started last
                night at a stroke while she
                was watching TV after eating;
                there 1 months she had had
                the same pain and it went
                alone
    Changes/evolutions: the pain is still
                present but a little less
                strong than at the beginning
    Location : at the stomach but goes to the
               right shoulder
    Timing/chronologie: in the beginning the
               pain is strong as a stab, it
               is still present but less
               strong.
    Quantification/severity: 4/10 but at the
               beginning it was stronger
Lifestyle:
    Addictions: she does not smoke and does
               not drink
    Pets      : no
Treatments:
    Type of treatment: no special treatment
    Method    : no
    Date and period: no
    Observation: no
```

Table 2: Clinical record of our standardized virtual patient.

the different dialogues (“Are you out of breath?”, “Do you have any shortness of breath?” for example);

variation of medical terms that refer to the same concept in the medical record to match more specialized terms to terms used in a general context (“Do you have abdominal pain?”, “Do you suffer from stomach pain?” or “Do you have high blood pressure?”, “Do you suffer from tension problems?” for example);

Category	Tag	Id
Aim of Consultation	MTF	00
Personal Data	PSN	01
Medical History	ATD	04
Symptoms	HDM	05
Lifestyle	MDV	02
Treatments	TRT	03
Other	UNK	06

Table 3: Tags used for annotation.

```
HDM0013 +++ ça fait mal à peu près où?
HDM0016 +++ c’est une douleur continue?
TRT0019 +++ êtes-vous sous traitements?
TRT0022 +++ vous ne prenez pas l’aspirine ?
ATD0025 +++ vous avez des antécédents médi-
            caux ?
```

Table 4: Some examples of annotated questions taken from the corpus.

the use of memory which aims to allow, for example in a consultation, the patient to refer to previous interactions to answer a question of the doctor (see an example in Table 8);

the use of ellipsis and anaphora as conversational markers illustrated with the words referring to or replacing the words used earlier in a dialogue (see an example in Table 8).

Table 8 lists some examples taken from the corpus to illustrate the different language phenomena.
It should be noted that our *Context QA dataset* is built for question answering systems that rely on the context in a dialogue to provide the response to a question.

```
P0014 +++  bah ça fait mal là, là, en haut du
           ventre, juste sous le sternum
P0017 +++  bah c’est une douleur qui est venue
           d’un seul coup comme ça, qui depuis
           n’a pas complètement disparu.
P0020 +++  non, je ne prends pas trop de médi-
           caments.
P0023 +++  non.
P0026 +++  c’est la première fois.
```

Table 5: Some examples of annotated responses taken from the corpus.

```
HDM0013 +++ P0014
HDM0016 +++ P0017
TRT0019 +++ P0020
TRT0022 +++ P0023
ATD0025 +++ P0026
```

Table 6: Cross between annotated questions and responses taken from the corpus.

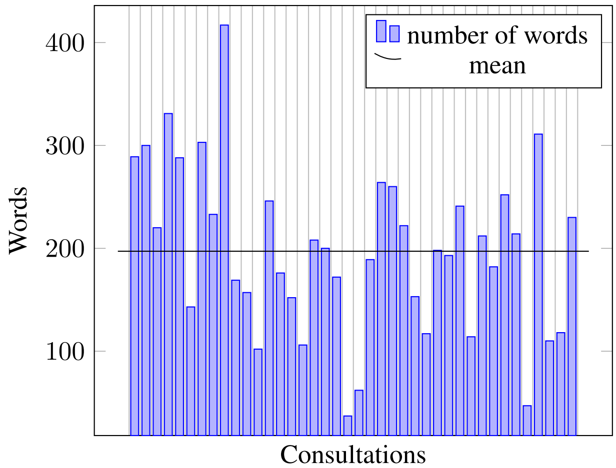

Figure 1: Number of words by consultations.

Tables 9, 10 and 11 list some global statistics of the current state of the corpus. These statistics include the number of consultations and sentences and the vocabulary size computed on all datasets. Tables 10 and 11 list the statistics by question types in order to show the proportion of sentences and words pronounced by the interns during the various consultations. They specify the contribution of the doctor and patient vocabulary by category of information present in the patient record. The total number and the average of sentences in a category are reported in Table 10 and the average number of words by category and by speaker in Table 11. The set of sentences includes both the questions and the responses while the calculation of the number of words considers the speeches of the doctor and the patient separately.
We also calculated the linguistic variables used in (Tanguy et al., 2011) to present the variations in the intern-VSP interactions during a consultation. The goal is to characterize the doctor-patient interactions contained in our *Context QA dataset*. The linguistic features considered are:

- the number of words of each consultation (N_{words});
- the number of words uttered by each speaker relative to the total number of words in the consultation ($speech_d$ for the doctor and $speech_p$ for the patient)
- the vocabulary specific to the doctor (voc_d) or the patient (voc_p) and the vocabulary common voc_comm to both in a consultation.

Figure 1 shows the variation of N_{words} with the average of the words pronounced by the doctor and the patient on the whole of the consultations.
In Figure 2, we show the evolution of the medical interactions according to the distribution of the speech of the

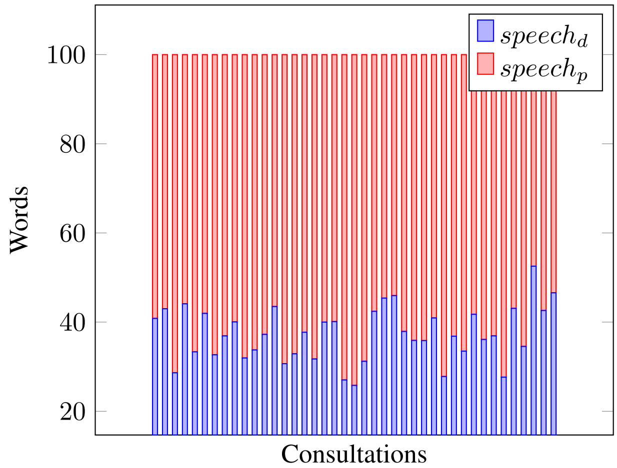

Figure 2: Number of words uttered by each speaker.

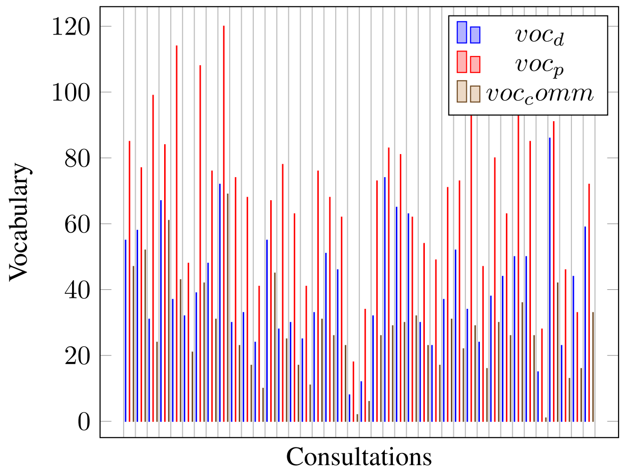

Figure 3: Vocabulary specific to each speaker and the common vocabulary.

doctor and the patient by plotting the variables $speech_d$ and $speech_p$. These rates correspond to the proportion of words uttered by each of them. The proportions show that the doctor has more talk than the patient except in a dialogue where both had a proportion of almost equal. By deploying the virtual patient to the medical school, the initial goal was to collect more data from the medical students. This can be noticed in our *Context QA dataset* by the high proportion of vocabulary and words specific to the doctor.
Figure 3 shows the proportion of vocabulary specific (voc_d and voc_p) and common (voc_comm) to both speakers. The common vocabulary consists of lemmas of nouns, verbs and adjectives that appear both in the doctor questions and the patient responses.

5. Medical Question Categorization Task

Corpora like this one are initially built in order to allow downstream applications to work. With the collected data, we conducted a doctor's questions categorization task to evaluate the benefits of our corpus. This task can be useful in user intention detection for the question and answer systems. We first enriched all of the data collected by using the concepts (classes of synonyms) created during the writ-

Id	Text	Description of Id				
		speaker	**dialogue**	**category**	**information type**	**speech turn**
d27000000	Bonjour	doctor	27	00	000	0
p27000000	Bonjour docteur	patient	27	00	000	0
d27000001	comment allez-vous ?	doctor	27	00	000	1
p27000001	J'ai mal au ventre	patient	27	00	000	1
d27050002	depuis quand ?	doctor	27	05	000	2
p27050002	Ça a commencé hier soir	patient	27	05	000	2
d27050083	et qu'est-ce que vous avez mangé?	doctor	27	05	008	3
p27050083	Une pizza	patient	27	05	008	3
d270100510	Quelle est votre profession?	doctor	27	01	005	10
d270100510	Je suis assistance maternelle	patient	27	01	005	10

Table 7: Example of annotated dialog extracted from the corpus.

```
structural variation
   doctor : Pourquoi êtes-vous venus aux
            urgences ?
   patient: J'ai très mal au ventre
   doctor : Qu'est-ce qui vous amène ici?
   patient: J'ai très mal au ventre
variation of medical terms
   doctor : Avez-vous de la fièvre?
   patient: je me sens fiévreux docteur
   doctor : Avez-vous de la température?
   patient: je me sens fiévreux docteur
memory
   doctor : Avez-vous pris des médicaments?
   patient: oui deux dolipranes
   doctor : combien en avez-vous pris?
   patient: deux
ellipsis
   doctor : Où avez-vous mal?
   patient: j'ai mal au ventre
   doctor : Depuis quand ?
   patient: ça a commencé hier soir
anaphora
   doctor : Comment est la douleur?
   patient: c'est comme un coup de
            poignard
   doctor : Elle est apparrue quand?
   patient: hier soir
```

Table 8: Some examples taken from the corpus.

	total
single-turn dataset	
#consultations	1
#sentences	5402
#vocabulary	2733
context QA dataset	
#consultations	41
#sentences	1818
#vocabulary	812

Table 9: Corpus statistics.

ing of the dialogue rules for the virtual patient for increasing the questions. This enriched data was then used to train and evaluate two classification models. The first one uses Convolution Neural Networks with an architecture similar to (Kim, 2014) and a linear classifier. The second model is based on FastText continuous word representation with rank constraints (Joulin et al., 2017). To train the different models, we represented the questions in a sequence of words used as inputs. We used pre-trained word embeddings (Bojanowski et al., 2017) for words from all datasets. The length of each sentence in the corpus is set to 50. This length is the average of the words in a question asked by the doctor in a dialogue. We used 10-fold cross-validation for training and validation with a ratio $90/10$ for splitting the corpus. The best-performing parameters that are used for the training and the evaluation of each model type are reported in Table 12.

		average
single-turn dataset		
Aim of Consultation	240	4.4
Personal Data	394	7.2
Medical History	1186	21.9
Symptoms	2544	47.0
Lifestyle	518	9.5
Treatments	266	4.9
Other	254	4.7
context QA dataset		
Aim of Consultation	148	10.4
Personal Data	120	8.4
Medical History	338	23.8
Symptoms	938	66.2
Lifestyle	96	6.7
Treatments	70	4.9
Other	82	5.7

Table 10: Total and average number of sentences by question type.

Results obtained on the validation data are shown in table 13. A convergence stability to 99% is obtained for the classifiers of Lifestyle (MDV) and Other (UNK) categories, while about more than half of Aim of consultation (MTF), Personal Data (PSN) and Medical History (ATD) classifiers have an accuracy that varies between 97% and 99%. Those in the Symptoms (HDM) category did not exceed 98%. We note that each of the obtained accuracies is significantly representative of the whole training and validation dataset.

	#words	
	#doctor	#patient
single-turn dataset		
Aim of Consultation	6	6
Personal Data	5	4
Medical History	6	8
Symptoms	7	7
Lifestyle	5	5
Treatments	7	8
Other	6	6
context QA dataset		
Aim of Consultation	4	5
Personal Data	5	4
Medical History	8	11
Symptoms	8	14
Lifestyle	5	5
Treatments	9	14
Other	5	5

Table 11: Average number of words by question type and speaker.

A more complete description of the system already built using this corpus and its performance is presented in (Laleye et al., 2020).
The corpus and the usage details will be made available[4] under a Free/Libre Open Source Licence. We will also include the training and evaluation scripts used in question categorization task to allow the reproduction of the results obtained with the corpus.

6. Conclusion

We explored clinical simulation sessions between medical students and a virtual patient to build an annotated corpus of dialogues in French. We adopted a data annotation scheme that allowed to prepare the statistics on questions and responses and also to characterize the interactions by presenting each speaker's vocabulary and common vocabulary. One of the benefits of the proposed corpus is that it can be used for downstream applications such as question answering, dialog management of a virtual patient, medical question categorization, clinical reasoning analysis and medical information extraction. We also demonstrated its benefit by using the proposed corpus in a question categorization task with excellent accuracies obtained on validation data.
In the future, the corpus could be enriched with dialogues and data related to other medical cases. We will also explore the possibility to replace our own-made categories (general and detailed) by categories from a standardized resource like UMLS (Bodenreider, 2004).

7. Acknowledgements

We wish to acknowledge the Agence Nationale de la Recherche (French National Research Agency): Project LabForSims2 15-IDFN-0008 funded under the IDEFI-Numérique Call for Projects.

[4]Temporary location during review: `https://github.com/kleag/labforsims2-corpus`.

8. Bibliographical References

Bodenreider, O. (2004). The Unified Medical Language System (UMLS): integrating biomedical terminology. *Nucleic Acids Research*, 32(Database issue):D267–D270, January.

Bojanowski, P., Grave, E., Joulin, A., and Mikolov, T. (2017). Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146.

Campillos Llanos, L., Thomas, C., Bilinski, E., Zweigenbaum, P., and Rosset, S. (2019). Designing a virtual patient dialogue system based on terminology-rich resources: Challenges and evaluation. *Natural Language Engineering*, pages 1–38, 07.

Celikyilmaz, A., Deng, L., and Hakkani-Tür, D., (2018). *Deep Learning in Spoken and Text-Based Dialog Systems*, pages 49–78. Springer Singapore, Singapore.

Gokcen, A., Jaffe, E., Erdmann, J., White, M., and Danforth, D. (2016). A corpus of word-aligned asked and anticipated questions in a virtual patient dialogue system. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 3174–3179, Portorož, Slovenia, May. European Language Resources Association (ELRA).

Harms, J., Kucherbaev, P., Bozzon, A., and Houben, G. (2019). Approaches for dialog management in conversational agents. *IEEE Internet Computing*, 23(2):13–22, March.

Iyyer, M., Yih, W.-t., and Chang, M.-W. (2017). Search-based neural structured learning for sequential question answering. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1821–1831, Vancouver, Canada, July. Association for Computational Linguistics.

Jin, L., White, M., Jaffe, E., Zimmerman, L., and Danforth, D. (2017). Combining cnns and pattern matching for question interpretation in a virtual patient dialogue system. In *Proceedings of the 12th Workshop on Innovative Use of NLP for Building Educational Applications*, pages 11–21, Copenhagen, Denmark, 01. Association for Computational Linguistics.

Joulin, A., Grave, E., Bojanowski, P., and Mikolov, T. (2017). Bag of tricks for efficient text classification. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, volume 2, pages 427–431, Valencia, Spain, 01. Association for Computational Linguistics.

Kim, Y. (2014). Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, 08.

Laleye, F. A. A., Blanié, A., Brouquet, A., Benhamou, D., and de Chalendar, G. (2020). Semantic similarity to improve question understanding in a virtual patient. In *TO APPEAR in Proceedings of the 35th ACM/SIGAPP Symposium On Applied Computing*, April.

Kernel size	**Kernel number**	**Embedding dimension**	**Dropout**	**learning rate**	**Optimizer**
3 to 5	300	300	0.4	3.0	Adamdelta

Table 12: Hyperparameter values.

	CNN	FastText
Aim of Consultation	98%	98%
Personal Data	98%	99%
Medical History	97%	99%
Symptoms	97%	98%
Lifestyle	99%	99%
Treatments	98%	98%
Unknown	99%	99%

Table 13: Accuracies per category on the validation data.

Laroche, R., Putois, G., and Bretier, P. (2010). Optimising a handcrafted dialogue system design. In *Proceedings of the 11th Annual Conference of the International Speech Communication Association (Interspeech)*, August.

Maicher, K., Danforth, D., Price, A., Zimmerman, L., Wilcox, B., Liston, B., Cronau, H., Belknap, L., Ledford, C., Way, D., Post, D., Macerollo, A., and Rizer, M. (2016). Developing a conversational virtual standardized patient to enable students to practice history-taking skills. *Simulation in Healthcare: The Journal of the Society for Simulation in Healthcare*, 12:1, 12.

Maicher, K., Zimmerman, L., Wilcox, B., Liston, B., Cronau, H., Macerollo, A., Jin, L., Jaffe, E., White, M., Fosler-Lussier, E., Schuler, W., Way, D., and Danforth, D. (2019). Using virtual standardized patients to accurately assess information gathering skills in medical students. *Medical Teacher*, 41:1–7, 06.

Porhet, C., Ochs, M., Saubesty, J., MONTCHEUIL, G. D., and Bertrand, R. (2017). Mining a Multimodal Corpus of Doctor's Training for Virtual Patient's Feedbacks. In *19th International Conference on Multimodal Interaction (ICMI)*, ICMI 2017- Proceedings of the 19th ACM International Conference on Multimodal Interaction, Glasgow, United Kingdom, November.

Spänig, S., Emberger-Klein, A., Sowa, J., Canbay, A., Menrad, K., and Heider, D. (2019). The virtual doctor: An interactive artificial intelligence based on deep learning for non-invasive prediction of diabetes. *CoRR*, abs/1903.12069:1–16.

Tanguy, L., Fabre, C., and Ho-Dac, L.-M. (2011). Caractérisation des échanges entre patients et médecins: approche outillée d'un corpus de consultations médicales. *Corpus*, pages 137–154, 01.