French FastContext: a Publicly Accessible System for Detecting Negation, Temporality and Experiencer in French Clinical Notes

Mehdi Mirzapour, Amine Abdaoui, Andon Tchechmedjiev, William Digan,
Sandra Bringay, Clement Jonquet

HAL Id: hal-03184001
https://imt-mines-ales.hal.science/hal-03184001

Submitted on 25 Apr 2022

French FastContext: A publicly accessible system for detecting negation, temporality and experiencer in French clinical notes

Mehdi Mirzapour [a], Amine Abdaoui [a], Andon Tchechmedjiev [a,b], William Digan [c], Sandra Bringay [a,d], Clement Jonquet [a,*]

[a] *Laboratory of Informatics, Robotics and Microelectronics of Montpellier (LIRMM), University of Montpellier, CNRS, France*
[b] *EuroMov Digital Health in Motion, Univ Montpellier, IMT Mines Alès, France*
[c] *Biomedical Informatics and Public Health Department, University European Hospital Georges Pompidou (HEGP), Paris, France*
[d] *University Paul-Valéry, Montpellier 3, France*

ABSTRACT

The context of medical conditions is an important feature to consider when processing clinical narratives. NegEx and its extension ConText became the most well-known rule-based systems that allow determining whether a medical condition is negated, historical or experienced by someone other than the patient in English clinical text. In this paper, we present a French adaptation and enrichment of FastContext which is the most recent, n-trie engine-based implementation of the ConText algorithm. We compiled an extensive list of French lexical cues by automatic and manual translation and enrichment. To evaluate French FastContext, we manually annotated the context of medical conditions present in two types of clinical narratives: (i) death certificates and (ii) electronic health records. Results show good performance across different context values on both types of clinical notes (on average 0.93 and 0.86 F1, respectively). Furthermore, French FastContext outperforms previously reported French systems for negation detection when compared on the same datasets and it is the first implementation of contextual temporality and experiencer identification reported for French. Finally, French FastContext has been implemented within the SIFR Annotator: a publicly accessible Web service to annotate French biomedical text data (http://bioportal.lirmm.fr/annotator). To our knowledge, this is the first implementation of a Web-based ConText-like system in a publicly accessible platform allowing non-natural-language-processing experts to both annotate and contextualize medical conditions in clinical notes.

1. Introduction and background

Health organizations store different kinds of medical notes within Electronic Health Records (EHRs). These notes often include unstructured elements (free text) that represent valuable information for medical research [20]. Therefore, the natural-language-processing community has developed automatic systems to detect clinical conditions and extract valuable knowledge from EHRs to facilitate: decision support [37], identification of patient cohorts [42], surveillance [21], etc. However, in order to extract the best outputs (selection of patients, prediction, decision) from EHRs, it is important to determine the context of the annotated clinical conditions. Indeed, it is not enough to detect a particular disease within the text, but a system shall also distinguish between a negated and an affirmed occurrence. Similarly, a patient might be excluded from a cohort if its record exhibits a condition that does actually concern a relative (e.g., parents). For this reason, researchers have proposed several methods to detect the context of –already annotated– clinical conditions, especially for English medical texts [7,22,40,19]. In this work, we have two objectives: (i) to propose an efficient equivalent system for French clinical text data (ii) to make it easily accessible and usable by the French medical community without any natural-language-processing expertise.

NegEx was, in 2001, one of the first reported systems to detect the context of clinical conditions [7]. It was originally developed to identify negated conditions in English discharge summaries and radiology reports. It uses a simple regular expression algorithm based on lexical cues (*trigger terms, pseudo-trigger terms, and termination terms*).[1] The system by default considers a condition *affirmed*, and marks it *negated* if it appears

---

[1] The ConText/FastContext algorithms and their different types of lexical cues will be presented in Section 2.1.

under the scope of a modifier (e.g., "no sign of", "absence of", "is ruled out"). Besides its simplicity, this algorithm is fast and effective. Subsequent research adopted two main types of negation detection methods for health narratives: Methods in the first category followed NegEx's idea using lexical cues [14], developed syntactic techniques based on part-of-speech tags and dependency trees [11], or combined both lexical cues and manually constructed grammar rules [22]. Methods in the second category used machine learning techniques such as Decision Trees [40], Conditional Random Fields or Support Vector Machines [32]. Machine Learning methods need annotated training data, which are usually difficult to get in clinical or medical environments. Furthermore, a comparative study has evaluated NegEx and two machine learning-based systems on the same dataset and concluded NegEx was more effective [17]. Wu et al reported that both categories are highly dependent on the clinical texts being evaluated [51]. Therefore, one may often prefer rule-based systems since they may be easily adapted to different corpora.

Due to its simplicity and effectiveness, NegEx has been adapted to at least four other European languages: Swedish [43], Spanish [9], French [13] and German [8,10]. Porting NegEx to these languages consisted mainly in translating its English lexical cues. In 2009, NegEx's successful approach was extended to detect temporality and experiencer of clinical conditions (always in English texts) with the development of the ConText [19] system. It offers new lexical cues to detect whether a condition is historical, hypothetical or experienced by someone other than the patient. At the same time, the authors improved negation detection with new lexical cues. The ConText system has been also ported to Swedish [48] and Dutch [3] but not to French. Moreover, NegEx/ConText has sometime been connected to English annotation systems so that they may both automatically identify the conditions and their contexts e.g., [6], however, we do not know any such integration for the French language, especially within a publicly accessible platform.

Regarding the French language, the previous NegEx adaptation obtained performance results comparable to the English version [13]. The authors' list of French lexical cues is available to the community, we therefore used it as a baseline in our development. However, this list contains a limited number of terms (270) and indeed, concerned only one context–negation–and did not cover temporality and experiencer. More recently, Garcelon et al. proposed another system to detect the context of clinical conditions in their systematic indexing/search of the Paris Necker Hospital's data warehouse [16]. In addition to negation, the authors compiled French lexical cues to detect family-related conditions. Their system is quite specific as it filters out negated and family-related conditions only in the context of the hospital's full-text search engine. Still, we used their released lexical cues to enrich our systems. A first version of our system, called French ConText, porting the entire ConText algorithm to French was released at the end of 2017 [2]; however this was an experimental prototype which was relying on an older (now deprecated) and less performant implementation. Beyond approaches to French context detection based on lexical cues, there is a significant body of work of context detection with more sophisticated linguistic approaches or supervised machine learning. Although cue-based approaches are very efficient for morphologically poor languages such as English, their ability to generalize their detection to morphological variations in more morphologically diverse languages like French is poorer (short of enumerating all possible variations of trigger terms). Dalloux et al. [12] propose a negation detection system based on a supervised approach (BiLSTM), which offers joint entity and negation detection performance of up to 76%.

In 2018, FastContext [41], a new implementation of ConText, was introduced. This implementation benefits from revised n-trie structures that allows fast execution of lexicon-based systems. FastContext has outperformed previous implementations of ConText. The n-trie-based engine of FastContext, implemented in Java,[2] consistently leads to gains in accuracy as the rule-set increases in size. Another important aspect of FastContext is its explicit externalized rules feature (i.e., the rules can be defined in a file separate from the implementation) that allows faster adaptation and reusability.

Regarding semantic annotation workflows to annotate French biomedical data with terminologies or ontologies, there are no open and public solutions available, despite of some research tools developed by Rouen University Hospital and since transferred to private companies (e. g., [38]). In the context of the Semantic Indexing of French biomedical Resources (SIFR) project, our group develops the SIFR Annotator (http://bioportal.lirmm.fr/annotator) to address the lack of off-the-shelf openly and easily accessible semantic annotation systems for French [23,45,47]. This service was significantly enhanced and customized for French,[3] but is originally based on the National Center for Biomedical Ontology (NCBO) Annotator (http://bioportal.bioontology.org/annotat or) [24]: a web service allowing to use available biomedical ontologies for annotating datasets automatically. The annotator service processes raw textual descriptions, tags them with relevant biomedical ontology concepts and returns the annotations to the users in several formats such as JSON-LD, RDF or BRAT. The SIFR Annotator uses 30 standard terminologies and ontologies collected in the SIFR BioPortal, a local instantiation of the NCBO technology [36,50] dedicated to French.[4] With the objective of proposing a complete ConText-capable semantic annotation workflow for French clinical notes, we have integrated French FastContext within the SIFR Annotator workflow as will be explained later. In this paper, we will focus solely on evaluating the context detection performance on a set of pre-identified entity annotations and not the performance of named entity recognition of SIFR Annotator, which has been evaluated thoroughly [47].

The SIFR Annotator not only identifies entities but also performs entity linking (normalisation) by associating explicit ontology classes to entities. Dalloux et al. also offer an online demonstrator or their system [12], however documentation is lacking, and the examples provided are non-functional, contrarily to our production-ready system. There have been more complex systems for the fine grained detection of experiencer [27,39] and time expressions [44,25,18], however their scope, objectives and use-cases are much broader than the type of detection ConText performs: they identify precise time expressions, dates and temporal relations, while ConText provides a very rough classification. The main use-case for ConText annotation is clearly for document indexing and classification rather than more advanced tasks such as de-identification of clinical text. There are in fact clinical deployments of our system that are used alongside HeidelTime for complementary purposes [33].

The evaluation of French FastContext was challenging due to the lack of a French gold standard and to the difficulty of accessing clinical data in general. If a few gold standard datasets have been produced for negation detection in English [49] and Spanish [28], there are no similar corpora for the French language. The sentences annotated in [8] and [16] are of restricted access, inaccessible to us for privacy-related restrictions. Therefore, with the collaboration of the HEGP hospital and agreement of the INSERM CépiDC department, we manually annotated the context of thousands of medical conditions in French clinical narratives. Wu et al have proved that negation detection performance suffers when there is no in-domain development [51]: a system performing

---

[2] The java implementation of FastContext is available in https://github.com/jianlins/FastContext .

[3] A paper detailing the implementation and evaluation of the SIFR Annotator is currently under preparation. The reader can nevertheless refer to [23], [45].

[4] Since 2019, the generic NCBO BioPortal technology is branded as the OntoPortal technology and jointly developed by the OntoPortal Alliance (https://ontoportal.org), under the guidance of Stanford University.

well on a specific corpus may lead to poor results when applied to another corpus. Therefore, we decided to evaluate French FastContext on two different types of health narratives: (i) death certificates and (ii) clinical notes from EHRs.

The remainder of this paper presents and discusses the adaptation and enrichment of FastContext to the French language, the integration of French FastContext within the SIFR Annotator workflow and its evaluation on the chosen types of clinical narratives.

2. Materials and methods

In this section, we first present the original ConText algorithm. Then, we describe its adaptation to the French language and its integration within the SIFR Annotator. Finally, we present the preparation and annotation of the gold standard corpora.

2.1. The ConText algorithm and FastContext implementation

We have used FastContext [41], a state-of-the-art implementation of the ConText algorithm, which is based on an n-trie engine which runs two orders of magnitude faster and is far less sensitive to rule set size compared to other implementations [19]. FastContext supports externalized rules that can be modified explicitly without any need for further code-level changes. In the domain of clinical information extraction, we typically deal with three or four different types of contextual cues. The Java implementation of FastContext includes a comprehensive set of 1322 rules refined from the original ConText rule set, supplemented by rules contributed from other resources [41], as well as rules designed by the FastContext developers. These rules cover four types of contextual cues: negation, certainty, temporality and experiencer. For each type of context, the cues may be one of (default values are presented in italic):

- **Negation:** *affirmed* or negated;
- **Certainty:** *certain* or uncertain;
- **Temporality:** *recent*, historical or hypothetical;
- **Experiencer:** *patient* or non-patient.

By default, a medical condition is considered as affirmed, certain, recent and patient. However, the contextual values may change if the condition falls within the scope of a modifier (also called a trigger term). For instance, in the sentence "the patient has no sign of melanoma," the underlined condition will be negated by ConText since it follows the modifier "no sign of". For each non-default value, the ConText system maintains a separate list of modifiers. For instance, the term "denies" triggers the value negated, the term "history of" triggers the value historical, while the term "mother's" triggers the value non-patient. Velupillai et al. [48] proposed a version of ConText considering even more complex contextual values.[5] In this work, initially in 2017, we started from the original version of ConText [19] for its simplicity and availability of a Java implementation. In 2019, we moved to FastContext [41] as an implementation that has increased speed and accuracy.[6]

The scope of each modifier is pre-defined as either forward (e.g., "denies") or backward (e.g., "is ruled out"). This scope ends by default either at the end of the sentence (full stop) or by setting an explicit window size configuration i.e., the number of tokens that come after a concept. However, some terms can override this default termination and act as termination terms. For instance, the word "but" ends the scope of negation modifiers such as in the sentence: "no sign of melanoma but multiple common moles". Finally, some expressions contain trigger terms but do not modify the context. For instance, the expression "no increase" contains a negation trigger term which is "no" but does not act as a modifier. In order to avoid false positives, the expression "no increase" will be added to a list of pseudo-modifiers which will not change the default contextual values. In conclusion, the FastContext system is based on three categories of lexical cues:

- **Trigger terms:** trigger a non-default contextual value of conditions falling within their scope;
- **Pseudo-trigger terms:** contain modifiers and do not change the context;
- **Termination terms:** end the scope of modifiers.

Adapting FastContext to the French language consisted in compiling an equivalent French list of these lexical cues (as described Section 2.2) and adapting some of the regular expressions to French. Table 1 presents some examples of English lexical cues and their French counterparts as used in our French rule set: "term, scope/category, context(s)."

2.2. Compilation of French lexical cues

Our list of French lexical cues has been obtained through automatic translation, manual validation and enrichment with previously existing alternate French lists. We decided to combine both machine translation and human expertise in order to create a comprehensive list of French lexical cues. Indeed, if human expertise should lead to good precision, machine translation allows us to assist humans in finding new terms that they may not think of. From the beginning of the project, we have pursued two different strategies in creation of French lexical cues that we detail below, followed by a merging process at the end.

2.2.1. Machine translation
Phase I: In this phase, our plan was to translate 356 terms of the English version of ConText into French. English online machine translation services were used to translate all the terms. Firstly, six dictionary-based automatic translators were queried: Babla, Sensagent, CNRS-ISC, Colins, Linguee, Wordreference.[7] Only the translated cues returned by at least three translators were kept. A prior study evaluated these translators for creating a sentiment lexicon and showed that a majority vote gives high-quality translation [1]. However, dictionary-based translators fail to handle compound terms (with multiple words). Therefore, a second list of French lexical cues was compiled using Google Translate. When several translations are proposed by Google Translate, we select only the first one which is the one with the

Table 1
Examples of English lexical cues and their French counterparts.

	English	French
Trigger terms	"absence of , pre, neg"	"absence de , pre ,neg"
	"is ruled out , post, neg"	"a été écarté , post, neg"
	"history of , pre, hist"	"antécédents de , pre, hist"
	"mom , pre, exp"	"mère , pre, exp"
Pseudo-trigger terms	"without difficulty , pseudo, neg"	"sans difficulté , pseudo, neg"
	"no history of , pseudo, hist"	"pas d'antécédents de , pseudo, hist"
Termination terms	"although , termin, neg"	"cependant , termin, neg"
	"complains , termin, histexp"	"se plaint , termin, histexp"

highest likelihood of success. At the end of this step, we automatically obtained a first list of 809 French lexical cues. This list was used in the first version of our work, French ConText released in 2017 [2].

Phase II: In this phase, our plan was to translate the 1322 existing English terms available in the FastContext implementation. We started with a semi-automatic translation of all the English terms to French with help of Google Translate. This provided our first seed term translations. The number of tokens for each term varied between one to eight. A list of relevant synonyms was assigned for maximum 3 tokens of each term. We obtained the synonyms from different sources such as: (i) synonym relations in Rezo-JDM [26], a French lexical-semantic network; (ii) French WordNet [4]; (iii) other available online synonym dictionaries, namely: Cnrtl (http://cnrtl.fr), synonymo (http://synonymo.fr), Crisco (http://crisco.unicaen.fr) and DicSyn (www.dictionnaire-synonymes.com). We manually curated the translation and removed the unrelated synonyms from the initial automatic generated lists. This simple strategy increased the possible number of translated terms. For instance, as for English term "adequate to rule out" we can have different valid French translations such as "suffisant pour écarter", "suffisant pour rejeter" and "suffisant pour éliminer", "satisfaisant pour écarter", "satisfaisant pour rejeter" and "satisfaisant pour éliminer". For the purpose of simplicity, we designed a compact Context-Free Grammar (CFG) based representation for the generation of translation variants. For example:

– **S** => **Adj** pour **Verb**
– **Adj** => *suffisant* | *satisfaisant*
– **Verb** => *écarter* | *rejeter* | *éliminer*

In which S, Adj and Verb are variables and synonyms are terminals. We grouped relevant French rules together and for each group we assigned several CFGs. After this semi-automatic processing, we obtained 737 CFGs that generate 10,147 French lexical cues from the 1322 initial English terms.

2.2.2. Manual validation

All the automatic translations have been checked, corrected and enriched manually by at least one bilingual biomedical text mining expert (among the authors). For instance, the English term "is ruled out" has been translated automatically to the French term "est écarté". We validated this translation but also added a new one: "est éliminé". Manual curation also consisted in adding different inflected forms from the automatically obtained translations (feminine/masculine, singular/plural). For instance, the term "is ruled out" may be used either for a masculine or a feminine occurrence, which is not the case in French (masculine: "est écarté"/"sont écartés", feminine: "est écartée"/"sont écartées").

For translations alternatives generated from phase I, we implemented a dedicated web-based application to check (validate or delete) each automatic translation and generated alternative, as well as to add new translations or new corrected forms. At the end of this step, we automatically removed duplicates and obtained 574 manually validated French lexical cues.

For translations generated from phase II, we set-up a Jupyter Notebook to allow for the manual validation of generated translation alternatives –mainly deletion, and some additions or modifications of the CFGs. The CFG-based representation enabled the validator to perform quick corrections by directly modifying the rules and by applying changes to multiple rules simultaneously.[8] After the validation step, we were left with 402 validated CFGs that generate 8,981 French lexical cues.

2.2.3. Enrichment

After manual validation of French lexical cues in phase I, our list of cues was enriched with the two previously available French lists cited in introduction [13,16]. We first, automatically compared our list with the two previous lists. We found that most of the terms from these two lists were already present in our own (with the same category and context). We then manually verified all the remaining cues that we did not already have. This process allowed us to enrich our list with new terms (such as: "élimine") and new forms (such as the apostrophe: "aucun signe d'"). Our final phase I list contained 710 entries (negation: 470, possible: 134, experiencer: 72, historical: 42, hypothetical: 22).[9]

The outputs of the phase I enrichment were directly merged to the outputs of phase II. Only 229 lexical cues were left and not already in phase II's list. Because the format of the lists were not exactly the same (phase I was in ConText's format were as phase II was in FastContext's format) we had to:

● Transform the *possible* cues into *uncertain* cues;
● Include default scope for each new added cues.

The final version of the French lexical cues consists of 10,147 French lexical cues (negation: 7668, non-patient: 528, uncertain: 1234, historical: 628, conditional: 89) after removing duplicated terms in the merging phase. This final list is publicly available for reuse by the community at https://github.com/practikpharma/FrenchFastContext .[10]

2.3. Integration within the SIFR Annotator

The US National Center for Biomedical Ontology offers a web service for text annotation with ontology concept [24]. This service originally designed for biomedical data was not very suitable for clinical text annotation. In order to add new functionalities to the NCBO Annotator, we have designed a proxy architecture that enables seamless extensions of any NCBO-like annotator service by pre-processing of the input text and parameters, and post processing of the annotations [46]. This proxy architecture has enabled us to integrate new features to both the SIFR and NCBO Annotator including contextualization, scoring, new output formats and coarse-grained semantic annotation (with UMLS Semantic Groups). For contextualization of recognized concepts, we have integrated ConText [19] for English in the NCBO Annotator+, and French FastContext, presented here, in the SIFR Annotator [47].[11] The SIFR Annotator is to the best of our knowledge the only openly available web service enabling both recognition and contextualization of concepts from 30 medical terminologies and ontologies in text. The incorporation of FastContext into the NCBO Annotator for English clinical text, much expected by BioPortal's users, was a critical addition especially for hospitals installing the NCBO virtual appliance to locally process clinical text.[12]

Automatic named entity recognition or semantic annotation (entity linking) of text is a prerequisite to the identification of clinical context, as modifiers are bound to particular concept mentions. Since the performance of SIFR Annotator has been extensively evaluated [47] on currently available French clinical notes corpora i.e., Quaero [35] and

---

[8] The Jupyter Notebooks we used to semi-automate the generation of the French lexical cues for FrenchContext are publicly available and documented for reproduction here: https://github.com/practikpharma/FrenchFastContext-CuesGeneration

[9] For archiving, this list is publicly available inside the French ConText code (see file ConTextFrench.java) : https://github.com/practikpharma/FrenchContext

[10] The list is versioned to keep track of its evolution: the version presented in the paper and used for evaluation in Section 3 is v1.

[11] Respectively accessible at http://bioportal.lirmm.fr/ncbo_annotatorplus and http://bioportal.lirmm.fr/annotator.

[12] Our contribution is also currently being integrated by our partner at Stanford's BMIR developing the original NCBO Annotator inside the NCBO BioPortal (http://bioportal.bioontology.org/annotatorplus).

CéPIDC [34] (the latter in the context of the participation of SIFR Annotator in the CLEF eHealth 2017 challenge [45])[13]. In this paper, we focus on the evaluation of the context identification performance of FastContext alone (P/R/F1) on concepts already identified by SIFR Annotator.

The SIFR Annotator workflow is composed of several steps: dictionary creation from ontologies/terminologies, concept recognition, semantic expansion (with concept-to-concept mappings and is-a hierarchy), and post-treatment of annotations. Since FastContext needs pre-annotated conditions, the contextual features have been added in the post-treatment step. Once the user text has been annotated, the new component splits the text into several sentences using the full stop, then applies French FastContext to each annotation with the sentence that contains this annotation as a parameter. Finally, the obtained contextual features are added to each annotation and formatted in the user's requested outputs.

Fig. 1 shows the SIFR Annotator user interface and the results obtained on a concrete example. The input text is written in French, it can be translated as the following: "the patient has no sign of melanoma, even if his father had a history of skin cancer". The results table shows the first annotation (melanoma) is negated, recent, and concerns the patient, while the second one (skin cancer) is affirmed, historical, and does not concern the patient. To reproduce this example with the web service use the following REST call:[14]

```
http://services.bioportal.lirmm.fr/annotator/?
text = Le%20patient%20n%27a%20aucun%20signe%20de%20 m%C3%A9lanome,
%20bien%20que%20son%20p%C3%A8re%20a%20des%20ant%C3%A9c%C3%
A9dents%20de%20cancer%20de%20la%20peau.
&ontologies = MSHFRE
&longest_only = true
&exclude_numbers = false
&whole_word_only = true
&exclude_synonyms = false
&expand_mappings = false
&fast_context = true
&score_threshold = 0
&confidence_threshold = 0
&lemmatize = false
&semantic_groups = DISO
&display_links = false&display_context = false
&apikey = 1de0a270-29c5-4dda-b043-7c3580628cd5
```

The SIFR Annotator user interface is mostly used as a demonstrator or by novice users. Indeed, the service is conceived to be called automatically using a web service application programming interface and can be plugged within data curation and annotations workflows. In addition, because access to clinical data is often restricted to hospital information systems only, we have designed a Docker-based implementation (https://github.com/sifrproject/docker-compose-bioportal) of the whole system that can be easily deployed locally; such as when used by the University European Hospital Georges Pompidou (HEGP).

2.4. Evaluation method

In order to assess the quality of French FastContext, we compiled a gold standard that associates each pre-tagged medical condition with its manually annotated contextual value.

2.4.1. Data sources

Since the type of the evaluation corpus has a huge influence of the context detection performance [51], we decided to evaluate French FastContext on two different types of health narratives:

– **Death certificates** obtained from the CépiDC causes of death corpus made available to the participants of CLEF eHealth 2017 task 1 [34]. This dataset contains 31,690 causes of death certificates as free-text descriptions, in French, reported by physicians using ICD-10 codes. Each document describes the death of only one person and is composed of a list of short sentences. On average, each sentence contains a little more than 8 words.
– **Electronic health records (EHRs)** obtained from the HEGP hospital describing patients' conditions and clinical reasoning. Data were confined inside the hospital and accessed only by authorized experts. The data in question contains around 4 million documents of various medical sub-fields. The large number of documents allowed us to extract more candidate sentences compared to the death certificates, especially for the experiencer evaluation. The sentences extracted from EHRs were also longer than those extracted from the death certificates. On average, each sentence contains a little more than 16 words.

2.4.2. Data preparation

For each data source, we semi-automatically created three test sets to evaluate each of the considered contexts (negation, certainty, temporality and experiencer). A first exploration of the data showed that most of the medical conditions were affirmed, recent and concerned the patient. Finding conditions with non-default contextual values was difficult due to their scarcity, which made the annotation time consuming. In order to augment the number of conditions with non-default contextual values, we adopted the following method:

– First, we selected candidate sentences[15] containing trigger terms and the same number of sentences not containing any trigger terms. This process allows us to improve the chance of finding negated, uncertain, historical, hypothetical and non-patient related conditions in the candidate sentences, while still being able to find new forms that were not included in our trigger terms list. The authors of the Swedish adaptation of NegEx [43] reported a similar method.
– Then, we used the SIFR Annotator to automatically annotate medical conditions in the candidate sentences. We restricted the annotation to keep only conditions tagged with a UMLS Semantic Type from the Semantic Group "Disorders (DISO)" [29] but we did not use any specific ontologies. For simplicity reasons, we kept only one condition per sentence –the first one appearing in the sentence–[16] and removed sentences without any annotated condition. The possible impact of this choice on the evaluation results are discussed Section 4.

2.4.3. Manual annotation

We manually annotated the context of each pre-tagged medical condition using the BRAT annotation tool (http://brat.nlplab.org). Each corpus was annotated by three human annotators. It is important to mention that the manual annotation only concerned the context of the automatically pre-tagged clinical conditions. Similarly, since the aim is only to evaluate French FastContext, the lexical cues that change the context of the medical conditions ('no sign of', 'history of', etc.) have not

---

[13] During the CLEF eHealth 2017 campaign, the SIFR Annotator obtained median results compared to the rest of the competitors; ahead of other knowledge-based systems but behind specifically tailored supervised learning systems. It is encouraging especially considering that we have not customized in any way the service to process these specific data.

[14] This web service call can be used "as this" for demo purposes. But for additional use of the SIFR Annotator we request users to register and use their own APIkey. See documentation here: http://data.bioportal.lirmm.fr/documentation and in [47].

---

[15] The sentence segmentation was performed simply using full stops or newlines.

[16] The first that has been detected by the SIFR Annotator which is not necessarily the real first appearing condition in the sentence.
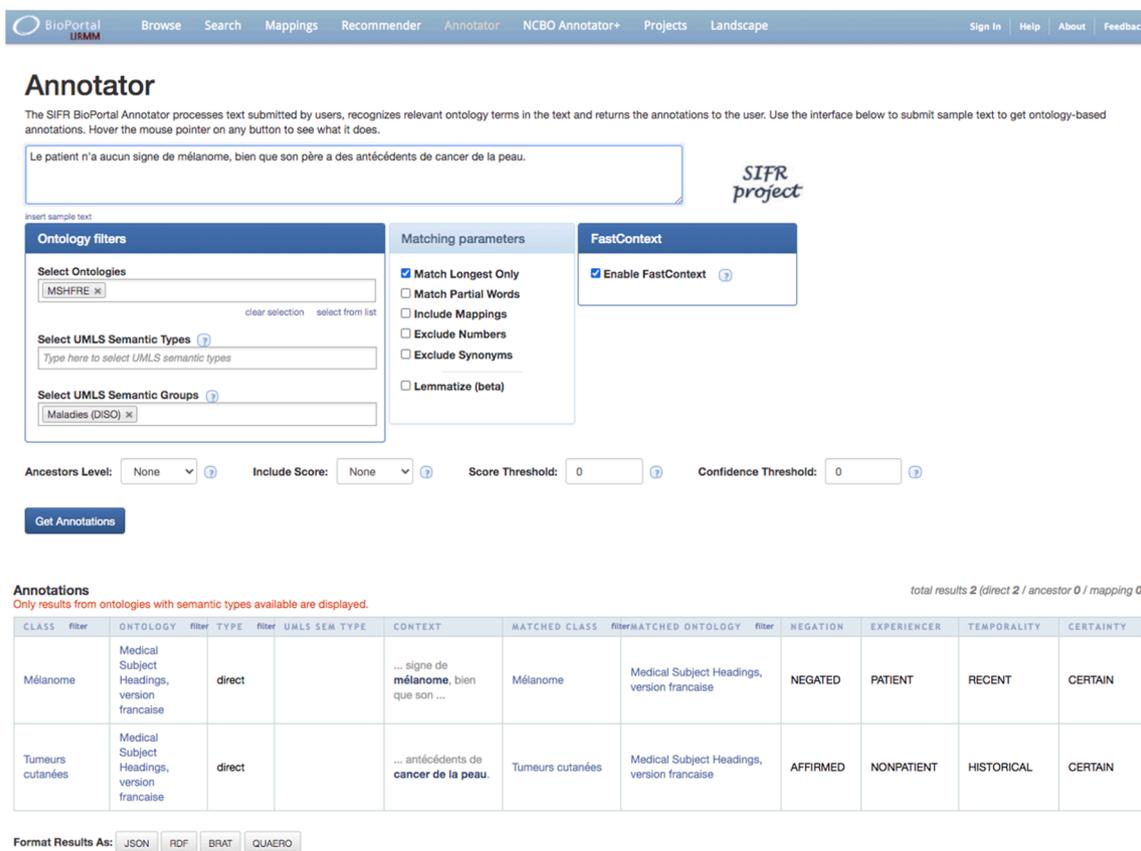
Fig. 1. SIFR Annotator user interface integrating the new contextual features. In this example, the target terminology is the French version of the Medical Subject Headings (MSHFRE) and the annotations are filtered to the 'Disorders' semantic group [29]. The semantic expansion (hierarchy and mapping) is not activated and only longest matches are kept.

been annotated either. For more details about the annotation process, please visit the annotation guidelines that have been made publicly available.[17] The whole process (data preparation and annotation) took a couple of weeks in order to compile a gold standard for both death certificates and EHRs. Additionally, the manual annotations of EHR data had to be done in-house on HEGP's premises for data access security and privacy reasons.

The annotators were asked to choose the appropriate contextual value for each pre-tagged medical condition. However, they observed that *possible* and *hypothetical* conditions are very rare in the annotated clinical narratives. Indeed, only 6 conditions were annotated as *possible*, while zero conditions were annotated as *hypothetical*. We decided to perform a bi-class evaluation only for each context (negated/affirmed, historical/recent, non-patient/patient) as previously conducted when evaluating NegEx/ConText-like systems [43,13]. Therefore, due to corpus limitations, we do not evaluate the performance of French FastContext in detecting *uncertain* and *hypothetical* conditions, even though the system is implemented to detect them.

Table 2 presents the number of manually contextualized sentences within each test set. Each sentence contains only one medical condition detected automatically by the SIFR Annotator. We also present the Fleiss kappa [15] agreement observed for each dataset.

3. Results

Using the manually annotated test sets presented above as gold standards, we have evaluated French FastContext. In this section, we

---

Table 2

Number and repartition of the annotated sentences between the different contextual values, as well as the average number of words in each sentence. The number of sentences is equal to the number of conditions.

		Negated	Affirmed	Total	Avg. # words	Kappa
Negation	Death certificates	40	990	1030	8	0.87
	Electronic health records	145	855	1000	17	0.85
		Historical	Recent	Total	Avg. # words	Kappa
Temporality	Death certificates	68	82	150	8	0.88
	Electronic health records	86	214	300	16	0.84
		Non-patient	Patient	Total	Avg. # words	Kappa
Experiencer	Death certificates	10	50	60	9	1
	Electronic health records	133	167	300	16	0.68

---

present results obtained in terms of precision, recall and F1-measure for each evaluated context (negation, temporality and experiencer).

3.1. Negation

Table 3 presents the results obtained by French FastContext and the previous adaptation to French of NegEx [13] on the same annotated sentences (obtained by using Déléger and al.'s trigger term list). French FastContext system obtains better F1-measures on both types of clinical narratives. Mainly, our system obtains better recall for the prediction of the negated class which means that it allows to find more negated conditions. This observation is clearly explained by the expanded list of lexical cues that we have compiled and used as described in Section 2.2. Furthermore, our results show slightly better performance than the English and Swedish versions of ConText one evaluated on different corpora. Indeed, these systems obtained between 0.80 and 0.82 F1-measures for the prediction of the negated class [43].

3.2. Temporality

Table 4 presents the results obtained by French FastContext on the annotated temporality sentences. Between 0.81 and 0.92 F1-measure is very good results for a first implementation for French. Indeed, the original English ConText algorithm obtained around 0.76 F1-measure for the prediction of the historical class [19].

3.3. Experiencer

Table 5 presents the results obtained by French FastContext on the annotated experiencer sentences. The reported F1-measures (between 0.79 and 0.98) can be considered good results for a first implementation for French. In comparison, the original English ConText algorithm obtained an F1-measure of 1.0 for the prediction of the non-patient class, but using a test set composed of only five conditions that have been annotated as experienced by someone other than the patient [19]. In our case, we have identified 360 conditions that are experienced by someone other than the patient.

4. Discussion

Here, we discuss the main limitations of this study, justify our choice of the ConText algorithm as a starting point, provide an error analysis of French FastContext implementation and present future extensions.

Evaluation: Even if the evaluation corpus considered two different types of clinical narratives, there are other specifications that may also influence the generalization of the presented results [51]. Hence, it is important to mention a few limitations regarding the evaluation corpus used in this work. First, our French FastContext system has been evaluated on only one type of medical entities (disorders). However, the SIFR Annotator is able to detect much more entities such as medications, procedures, anatomical components, etc. The choice of evaluating "disorders" has been motivated by the number of papers about clinical

Table 3
Comparative evaluation results between previous performance of French NegEx and our French FastContext system on both types of clinical narratives.

			Precision	Recall	F1
Death Certificates	French NegEx	Negated	1	0.675	0.806
		Affirmed	0.987	1	0.993
	French	Negated	0.923	0.900	0.911
	FastContext	Affirmed	0.996	0.997	0.996
Electronic Health Records	French NegEx	Negated	0.989	0.648	0.783
		Affirmed	0.901	0.999	0.947
	French	Negated	0.916	0.759	0.830
	FastContext	Affirmed	0.956	0.987	0.971

Table 4
Evaluation results of the temporality detection on both types of clinical narratives.

		Precision	Recall	F1
Death Certificates	Historical	0.898	0.841	0.869
	Recent	0.890	0.931	0.910
Electronic Health Records	Historical	0.986	0.692	0.814
	Recent	0.859	0.995	0.922

Table 5
Evaluation results of the experiencer detection on both types of clinical narratives.

		Precision	Recall	F1
Death Certificates	Patient	0.980	0.980	0.980
	Non-Patient	0.900	0.900	0.900
Electronic Health Records	Patient	0.761	0.927	0.836
	Non-Patient	0.905	0.705	0.792

context detection that considered this type of medical entities (alone or among others) in their evaluation corpora [19,51,8,48], including the original version of ConText that has been evaluated only for symptoms and disorders. In addition, detecting disorders and conditions was actually the driving use case of our project with HEGP hospital, that have motivated this work: extracting pharmacogenomics knowledge from French EHRs to compare them to state-of-the-art knowledge published in scientific articles and references databases [31]. Therefore, we have evaluated French FastContext for the most appropriate medical category (disorders) but significant additional work would be required to evaluate the applicability of our system to contextualize other types of medical entities e.g., drugs ("the patient is not treated by drug X yet").

Second, we have selected the candidate sentences in such a way as to have a balanced corpus where half of the selected candidate sentences should contain trigger terms and half of them should not. This willing selection made the class distribution of the evaluation corpus different from the class distribution in the original data –which might result in better performance measures. Indeed, the proportion of conditions having default contextual values (for example *affirmed*) is certainly higher in the original data. Therefore, the repartition of the annotated sentences between the different contextual values presented in Table 2 concerns the evaluation corpus and is not representative of the source corpora used to generate it. The results presented in Tables 3, 4 and 5 are also influenced by the corpus development choices, but they still consider sentences that do not contain any of our trigger terms. Moreover, as explained Section 2.4.2, the annotation protocol only considers the first annotated concept in each sentence for the FastContext evaluation, which may limit the relevance of the evaluation with regard to the quality of the termination terms: in complex sentences, we may miss errors that would occur on conditions mentioned later in the sentence due to a missing or incorrect termination term –but only when another condition explicitly recognized by the SIFR annotator would be in the sentence before. Plus, by picking up only one condition per sentence, our evaluation qualifies the capacity of the system to contextualize "conditions" and not mixed "condition sentences".

Finally, we could not find uncertain or hypothetical conditions in the corpora we used. Therefore, we could not evaluate French FastContext in detecting uncertain and hypothetical conditions. HEGP hospital's corpus of EHRs constantly grows with time as more data are anonymized and extracted from the medical information system [52]. We hope that a later evaluation would include sufficient amounts of uncertainty and hypothetical modifiers to enable this evaluation.

Lexical cues: Third, despite our efforts to come with a generic list of lexical cues as good as possible, our experiments have shown that customization of the list is often needed for a specific use case. For example, despite a few exceptions (e.g., 'mère' or 'mere' (for mother)), we have

not taken care of spelling mistakes in the list of cues. This will be a future improvement. For now, users will need to rely on an external spelling corrector when preprocessing the text to annotate. As another example, the scope of lexical cues might need to be customized too. We found a small change in the scopes of frequent lexical cues can have a significant impact on performance. We believe the externalization of the rules provided by FastContext and thus the SIFR Annotator (possibly locally installed) made such customizations easily doable.

Choice of ConText algorithm: Despite recently confirmed high performance obtained by state-of-the-art deep learning systems, in 2016, but still today, the choice of ConText to contextualize clinical annotations within the SIFR Annotator was also strongly justified by the fact that it can be easily plugged in the existing workflow. The per-sentence approach and the regular expression methodology of NegEx/ ConText/FastContext were quite relevant for the SIFR Annotator. The additional processing can be included in the post-treatment step in order to add contextual features to the concepts annotated in previous steps (either direct annotations or semantically expanded ones). For instance, the SIFR Annotator can generate a negated annotation with the concept 'cancer' (D009369) using MeSH hierarchy when the given text is "no sign of melanoma" because the concept 'melanoma' (D008545) has been directly identified in the text, then expanded with MeSH hierarchy and then negated. Another reason for that choice was the possibility to generalize the implementation done for the SIFR Annotator, to any NCBO-like annotator using a web service proxy architecture offering new functionalities (e.g., scoring [30]) by pre-processing of the input text and post-processing of the original results. Consequently, as presented in [46], the SIFR BioPortal now offers the possibility of querying the original English-text-based NCBO Annotator and contextualizes its annotations. The NCBO Annotator+ (http://bioportal.lirmm.fr/ncbo _annotatorplus) offers the same parameters, user interface and results but in the backend, relies on the original NCBO Annotator and its compiled dictionary from +800 English biomedical ontologies and terminologies.[18]

Error analysis: We evaluated French FastContext and compared it with another available system for negation detection thanks to manually annotated death certificates and EHRs. Experimental evaluation shows that French FastContext obtains better results on both types of clinical narratives due to its extended list of lexical cues. However, error analysis conducted by comparing the gold standard corpus and the French FastContext results allowed us to identify 72 new trigger terms that were not present in our list of lexical cues. Obviously, these new trigger terms were not integrated in the version of French FastContext evaluated here but were later included within the available web service. All these new terms have equivalents in English that are not present in the original FastContext list of lexical cues. For example, the term "non identifié" was not present in our list because its direct English translation "un-identified" was not present in the original FastContext list. Therefore, we may also enrich the original English version of FastContext by translating these new terms. However, we believe that adding more trigger terms should have a limited impact on the results (except for another kind of data). Indeed, it has been demonstrated that only few terms occur a large number of times while many terms occur a small number of times [8]. Finally, as all ConText-based systems, the performance of French FastContext depends on the quality of the text segmentation into sentences. In our evaluation, this segmentation was currently based on the presence of a full stop or a newline. However, we have observed this is not always the case in clinical narratives. Therefore, in the future, we are planning to use more developed sentence segmentation tools [5] and plug them in into the SIFR Annotator workflow.[19] We could also use the same segmentation tool used to evaluate the original ConText to make the comparison more accurate.

Superimposed errors: We are presenting here the evaluation method for the performance of FastContext alone, independently from the concept recognition performance of SIFR Annotator, since we want an evaluation that is comparable to the original evaluation of FastContext. Because errors compound geometrically on superimposed annotation tasks (i.e., first recognizing concepts and then contextualizing them) means that any true or false positive in the concept identification automatically results in a true or false positive in the context detection evaluation. Consequently, the joint recognition performance is roughly equal to the product of the performance ratios (i.e., in percentage points of P, R or F1). In our previous extensive evaluation efforts of SIFR Annotator for concept identification [47], the concept recognition performance varies between 50% F1 and almost 70% F1. For instance, if we then have a negation recognition performance of say 80% F1, we will get a combined performance of roughly 40% in the worst case and 56% in the best case. This fact must be kept in mind when assessing the results of the subsequent evaluation of French FastContext.

Extensions: In this paper, we restricted our evaluation to medical annotations considered as disorders. It will be interesting to evaluate the ability of French FastContext to contextualize other annotations and to see if the current list of lexical cues can be generalized to other types of biomedical annotations. On another hand, some annotations do not have to be contextualized. For instance, medical devices or anatomy parts. Plus, some terminologies natively include some trigger terms (e.g., "aucun signe de" in SNOMED International v3.5) which makes no sense to contextualize. Even if these cases may be easily ignored, in the future, we will implement a way to remove them. Finally, we have been working on the development of a new entity recognition module based on the Unitex toolkit (http://unitexgramlab.org/). We are considering implementing FastContext in the form of Unitex graph patterns that will allow us to capture more intricate and precise contextual clinical information.

5. Conclusion

In this paper, we presented an adaptation and enrichment of Fast-Context to the French language. The proposed system allows to detect *negation*, *temporality* and *experiencer* of already annotated clinical conditions (and technically speaking, *certainty* too). The system has been evaluated on two types of clinical narratives and obtained comparable results to the English and the Swedish versions of ConText. When compared on the same datasets, French FastContext outperforms the currently available adaptation of NegEx to the French language due its extended list of French lexical cues. Furthermore, our system has been integrated in the SIFR Annotator in order to make it easily and publicly accessible for the biomedical community. This service can be accessed manually using a web user interface and automatically using a web service API. In addition, the SIFR Annotator can be deployed locally to process sensitive data in-house (thanks to Docker packaging). Our implementation has also been generalized to English with the annotations produced by the NCBO Annotator. In both French and English, we believe we significantly enhanced the functionalities of the original annotators especially for annotating clinical data. Our team is collaborating with the HEGP Hospital in Paris and the French INSERM (National Institute of Health and Medical Research) at Nancy university hospital to exploit the SIFR Annotator in concrete medical research environments. In both cases, a local installation of the SIFR Annotator (including French FastContext) has been set up because of data access restrictions. In addition to medical annotation, our system will be able to filter out negated conditions and those experienced by someone other than the patient, track the patient history of medical events, therefore helping the medical experts analyze their clinical data.

---

[18] Note that in 2019, Stanford has integrated the NCBO Annotator + in the original NCBO BioPortal (http://bioportal.bioontology.org/annotatorplus).

[19] We are currently working to include a StanfordNLP developed sentence segmentation component.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Author Contributions

MM and AA compiled the trigger term list, drove the evaluation and wrote together most of the manuscript. AT offered expertise in natural language processing and helped with the evaluation and writing. AA and AT jointly implemented the system. MM created phase II CFG-based group of the rules. WD helped with the evaluation and access to EHR data. SB helped with the supervision and driving of the project. CJ conceived the project, provided the scientific directions and extensively edited this manuscript. All authors declare the absence of any conflicts of interest and approve the final manuscript.

References

[1] A. Abdaoui, J. Azé, S. Bringay, P. Poncelet, FEEL: a French Expanded Emotion Lexicon, Language Resources and Evaluation, p. 1–23, 2016.

[2] A. Abdaoui, A. Tchechmedjiev, W. Digan, S. Bringay, C. Jonquet. French ConText: Détecter la négation, la temporalité et le sujet dans les textes cliniques Français. SIIM: Symposium sur l'Ingénierie de l'Information Médicale, Toulouse, France, Nov 2017.

[3] Z. Afzal, E. Pons, N. Kang, M.C. Sturkenboom, M.J. Schuemie, J.A. Kors, ContextD: an algorithm to identify contextual properties of medical terms in a Dutch clinical corpus, BMC Bioinform. 15 (1) (2014) 373.

[4] S. Benoît, F. Darja, Building a free French wordnet from multilingual resources, in: Ontolex 2008 in association with 6th International Conference on Language Resources and Evaluation (LREC2008), Marrakech, Maroc, 2008.

[5] S. Bird, NLTK: the natural language toolkit, COLING/ACL on Interactive presentation sessions, Sydney, Australia, p. 69–72, july 2006.

[6] B.E. Chapman, S. Lee, H.P. Kang, W.W. Chapman, Document-level classification of CT pulmonary angiography reports based on an extension of the ConText algorithm, J. Biomed. Inform. 44(5) (oct. 2011) p. 728–737.

[7] W.W. Chapman, W. Bridewell, P. Hanbury, G.F. Cooper, B.G. Buchanan, A simple algorithm for identifying negated findings and diseases in discharge summaries, J. Biomed Inform. 34 (5) (2001) 301–310.

[8] W.W. Chapman, et al., Extending the NegEx lexicon for multiple languages, Stud. Health Technol. Inform. 192 (2013) 677–681.

[9] R. Costumero, F. López, C. Gonzalo-Martín, M. Millan, E. Menasalvas, An approach to detect negation on medical documents in Spanish, in: International Conference on Brain Informatics and Health, Springer, Cham, 2014, p. 366–375.

[10] V. Cotik, R. Roller, F. Xu, H. Uszkoreit, K. Budde, D. Schmidt, Negation detection in clinical reports written in German, in: Fifth Workshop on Building and Evaluating Resources for Biomedical Text Mining (BioTxtM2016), p. 115–124, 2016.

[11] V. Cotik, V. Stricker, J. Vivaldi, H. Rodriguez, Syntactic methods for negation detection in radiology reports in Spanish, in: 15th Workshop on Biomedical Natural Language Processing, BioNLP 2016: Berlin, Germany, August 12, 2016. Association for Computational Linguistics, 2016. p. 156-165.

[12] C. Dalloux, V. Claveau, N. Grabar, Détection de la négation: corpus français et apprentissage supervisé, Symposium sur l'Ingénierie de l'Information Médicale (SIIM 2017), Toulouse, France, 2017, p. 1–8.

[13] L. Deléger, C. Grouin, Detecting negation of medical problems in French clinical notes, in: 2nd ACM SIGHIT International Health Informatics Symposium, 2012, p. 697–702.

[14] P.L. Elkin, et al., A controlled trial of automated classification of negation from clinical notes, BMC Med. Inform. Decis. Mak. 5 (1) (2005) 13.

[15] J.L. Fleiss, Measuring nominal scale agreement among many raters, Psychol. Bull. 76 (5) (1971) 378.

[16] N. Garcelon, A. Neuraz, V. Benoit, R. Salomon, A. Burgun, Improving a full-text search engine: the importance of negation detection and family history context to identify cases in a biomedical data warehouse, J. Am. Med. Inform. Assoc. 24 (3) (2016) 607–613.

[17] S. Goryachev, M. Sordo, Q.T. Zeng, L. Ngo, Implementation and evaluation of four different methods of negation detection, Decision Systems Group, Boston, Maryland, DSG, Technical Report, 2006.

[18] T. Hamon, N. Grabar, Tuning HeidelTime for identifying time expressions in clinical texts in English and French, in: 5th International Workshop on Health Text Mining and Information Analysisin (EACL 2014 Workshop: Louhi), 2014, p. 101–105.

[19] H. Harkema, J.N. Dowling, T. Thornblade, W.W. Chapman, ConText: an algorithm for determining negation, experiencer, and temporal status from clinical reports, J. Biomed Inform. 42 (5) (2009) 839–851.

[20] K. Häyrinen, K. Saranto, P. Nykänen, Definition, structure, content, use and impacts of electronic health records: a review of the research literature, Int. J. Med. Inf. 77 (5) (2008) 291–304.

[21] V. Herasevich, et al., Limiting ventilator-induced lung injury through individual electronic medical record surveillance, Crit. Care Med. 39 (1) (2011) 34–39.

[22] Y. Huang, H.J. Lowe, A novel hybrid approach to automated negation detection in clinical radiology reports, J. Am. Med. Inform. Assoc. 14(3) (2007) p. 304–311.

[23] C. Jonquet, A. Annane, K. Bouarech, V. Emonet, S. Melzi, SIFR BioPortal: Un portail ouvert et générique d'ontologies et de terminologies biomédicales françaises au service de l'annotation sémantique, in: 16th Journées Francophones d'Informatique Médicale, (JFIM'16), Genève, Suisse, 2016, p. 16.

[24] C. Jonquet, N.H. Shah, M.A. Musen, The open biomedical annotator, Summit on translational bioinformatics, 2009, p. 56–60.

[25] R. Kessler, X. Tannier, C. Hagège, V. Moriceau, A. Bittar, Extraction de dates saillantes pour la construction de chronologies thématiques, TAL, vol. 53, p. 57-86, 2012.

[26] M. Lafourcade, Making people play for Lexical Acquisition with the JeuxDeMots prototype, in: 7th International Symposium on Natural Language Processing (SNLP 2007), Pattaya, Chonburi, Thailand, Dec 2007, pp.7.

[27] A. Lumbreras, Automatic role detection in online forums, Doctoral dissertation, Université de Lyon, 2016.

[28] M. Marimon, J. Vivaldi, N. Bel, Annotation of negation in the IULA Spanish Clinical Record Corpus, Workshop Computational Semantics Beyond Events and Roles, 2017, p. 43–52.

[29] A.T. McCray, A. Burgun, O. Bodenreider, Aggregating UMLS semantic types for reducing conceptual complexity, Stud. Health Technol. Inform., 84(0 1) (2001) p. 216–220.

[30] S. Melzi, C. Jonquet, Scoring semantic annotations returned by the NCBO Annotator, in: 7th International Semantic Web Applications and Tools for Life Sciences, SWAT4LS'14, Berlin, Germany, 2014, vol. 1320, p. 15.

[31] P. Monnin, J. Legrand, G. Husson, P. Monnin, J. Legrand, G. Husson, P. Ringot, A. Tchechmedjiev, C. Jonquet, A. Napoli, A. Coulet, PGxO and PGxLOD: a reconciliation of pharmacogenomic knowledge of various provenances, enabling further comparison, BMC Bioinform. 20 (2019) p.139.

[32] R. Morante, W. Daelemans, A metalearning approach to processing the scope of negation, in: 13th Conference on Computational Natural Language Learning, 2009, p. 21–29.

[33] V. Moriceau, X. Tannier, French resources for extraction and normalization of temporal expressions with heideltime, in: 9th International Conference on Language Resources and Evaluation (LREC14), 2014, p. 3239–3243.

[34] A. Névéol et al., A. Névéol, R.N. Anderson, K.B. Cohen, C. Grouin, T. Lavergne, G. Rey, A. Robert, C. Rondet, P. Zweigenbaum, CLEF eHealth 2017 Multilingual Information Extraction task overview: ICD10 coding of death certificates in English and French, in: CLEF 2017 Evaluation Labs and Workshop: Online Working Notes, CEUR Workshop Proceedings, 2017, p. 17.

[35] A. Névéol, C. Grouin, J. Leixa, S. Rosset, P. Zweigenbaum, The QUAERO French Medical Corpus: a ressource for medical entity recognition and normalization, in : Proc of BioTextMining Work, 2014, p. 24–30.

[36] N.F. Noy, et al., BioPortal: ontologies and integrated data resources at the click of a mouse, Nucleic Acids Res. mai 2009 37(web server), p. 170–173.

[37] P.J. O'Connor, et al., Impact of electronic health record clinical decision support on diabetes care: a randomized trial, Ann. Fam. Med. 9 (1) (2011) 12–21.

[38] S. Pereira, A. Névéol, G. Kerdelhué, E. Serrot, M. Joubert, S.J. Darmoni, Using multi-terminology indexing for the assignment of MeSH descriptors to health resources in a French online catalogue, in: American Medical Informatics Association Annual Symposium, AMIA'08, Washington DC, USA, 2008, p. 586–590.

[39] C. Pertin, C. Deccache, R. Gagnayre, T. Hamon, User profile detection in health online Fora, Stud. Health Technol. Inform. 247 (2018) 730–734.

[40] L. Rokach, R. Romano, O. Maimon, Negation recognition in medical narrative reports, Inf. Retr. 11 (6) (2008) 499–538.

[41] J. Shi, J.F. Hurdle, Trie-based rule processing for clinical NLP: A use-case study of n-Trie, making the ConText algorithm more efficient and scalable, J. Biomed. Inform., August 6, 2018.

[42] C. Shivade, et al., A review of approaches to identifying patient phenotype cohorts using electronic health records, J. Am. Med. Inform. Assoc. 21 (2) (2013) 221–230.

[43] M. Skeppstedt, Negation detection in Swedish clinical text: An adaption of NegEx to Swedish, J. Biomed. Semant. 2 (3) (2011) S3.

[44] M.D. Tapi Nzali, X. Tannier, A. Névéol, Automatic extraction of time expressions across domains in French narratives, in: Conference on Empirical Methods in Natural Language Processing (EMNLP 2015), Lisbonne, Portugal, 2015, p. 492–498.

[45] A. Tchechmedjiev, A. Abdaoui, V. Emonet, C. Jonquet, ICD10 coding of death certificates with the NCBO and SIFR Annotator(s) at CLEF eHealth 2017 Task 1, in: Working Notes of CLEF eHealth Evaluation Lab, Dublin, Ireland, 2017, vol. 1866, p. 16.

[46] A. Tchechmedjiev, A. Abdaoui, V. Emonet, S. Melzi, J. Jonnagaddala, C. Jonquet, Enhanced functionalities for annotating and indexing clinical text with the NCBO Annotator+, Bioinformatics 34 (11) (2018) 1962–1965.

[47] A. Tchechmedjiev, A. Abdaoui, V. Emonet, S. Zevio, C. Jonquet, SIFR annotator: ontology-based semantic annotation of French biomedical text and clinical notes, BMC Bioinform. 19(1), p. 405, nov. 2018.

[48] S. Velupillai, et al., Cue-based assertion classification for Swedish clinical text—Developing a lexicon for pyConTextSwe, Artif. Intell. Med. 61 (3) (2014) 137–144.

[49] V. Vincze, G. Szarvas, R. Farkas, G. Móra, J. Csirik, The BioScope corpus: biomedical texts annotated for uncertainty, negation and their scopes, BMC Bioinform. 9 (11) (2008) S9.

[50] P.L. Whetzel, N. Team, NCBO Technology: Powering semantically aware applications, Biomed. Semant., vol. 4S1, no S8, p. 49–58, avr. 2013.

[51] S. Wu, et al., Negation's not solved: generalizability versus optimizability in clinical natural language processing, PloS One 9 (11) (2014), e112774.

[52] A.-S. Jannot, A. Burgun, E. Thervet, N. Pallet, The diagnosis-wide landscape of hospital-acquired AKI, CJASN 12 (6) (June 2017) 874–884.