

Automatic Selection of Clinical Trials Based on A Semantic Web Approach

Marc Cuggia^{abg}, Boris Campillo-Gimenez^{acg}, Guillaume Bouzille^{abg}, Paolo Besana^{ag}, Wassim Jouini^{ag}, Jean-Charles Dufour^d, Oussama Zekri^c, Isabelle Gibaud^e, Cyril Garde^f, Regis Duvaufier^{bg}

^aINSERM, U1099, Rennes, F-35000, France, ^bCHU Rennes, CIC Inserm 1414, Rennes, F-35000, France

^cCRLCC Eugène Marquis, ^dLERTIM – Université de la méditerranée, ^eSIB, ^fENOVACOM, Université de Rennes 1, ^gLTSL, Rennes, F-35000, France

Abstract

Recruitment of patients in clinical trials is nowadays preoccupying, as the inclusion rate is particularly low. The main identified factors are the multiplicity of open clinical trials, the high number and complexity of eligibility criteria, and the additional workload that a systematic search of the clinical trials a patient could be enrolled in for a physician. The principal objective of the ASTEC project is to automate the prescreening phase during multidisciplinary meetings (MDM). This paper presents the evaluation of a computerized recruitment support systems (CRSS) based on semantic web approach. The evaluation of the system was based on data collected retrospectively from a 6 month period of MDM in Urology and on 4 clinical trials of prostate cancer. The classification performance of the ASTEC system had a precision of 21%, recall of 93%, and an error rate equal to 37%. Missing data was the main issue encountered. The system was designed to be both scalable to other clinical domains and usable during MDM process.

Keywords:

Clinical research, Decision support system, Semantic interoperability, automatic reasoning.

Introduction

Clinical trials (CTs) are the gold standard for testing therapies or new diagnosis techniques that would improve clinical care. CTs often rely on adequate sample sizes, but often remain incomplete or are cancelled due to missed recruitment targets within a certain timeframe and financial cost. The recruitment process faces many barriers that have already been well identified in the literature. The automation of the patient screening process by computerized recruitment support systems (CRSS) should address some of the recruitment barriers. In a review paper, we analysed the advantages and drawbacks of each CCRS described in the literature[1]. Based on this analysis, we developed a system based on the 3 following principles:

1. Electronic health records (EHRs) and eligibility criteria (EC) are usually written by humans for humans. Consequently, their formalization for an automatic screening system can prove to be challenging. An CRSS should rather process structured and coded data. Free text formulation of patient data and eligibility criteria results in too many ambiguities. Despite the efficiency of NLP methods, coded and structured data with a formal semantic remain the best situation to successfully develop and apply automatic reasoning methods.
2. To be scalable, a CRSS should be connected to any kind of clinical EHR source. Indeed, data from the same patient might be scattered in different hospital

information systems. To address this issue, an CRSS should adopt a service-oriented architecture and use semantic interoperability standards to communicate with the different data sources.

3. Similarly to any decision support system, CRSS should provide useful information about the eligibility status of patients strategically at the right time and place, such as when physicians decide on treatment plans. For instance in oncology, this decision could be taken during MultiDisciplinary Meetings (MDM).

ASTEC (Automatic Selection of clinical Trials based on Eligibility Criteria) is a French national research project that aims to develop an CRSS designed on the 3 principles mentioned above. This system has been tested by the Centre Eugene Marquis (CEM), a Regional Comprehensive Cancer Centre located at Rennes (Brittany, France). In this paper, we present the evaluation of the system.

Background

Multidisciplinary meeting and the recruitment process in clinical trials: MDM are an integrated team approach to health care in which medical and allied healthcare professionals consider all relevant treatment options and develop collaboratively an individual treatment plan for each patient. That is, MDM is about all relevant health professionals discussing options and making joint decisions about treatment and supportive care plans, taking into account the personal preferences of the patient. In France, multidisciplinary decisions are mandatory for all oncology patients.

Decision support systems for recruitment: For over 25 years, many attempts have been made to develop methods and tools supporting the recruitment process. More recently, projects with new technologies recent are arising, such as the Biomedical Translational Research Information System (BTRIS) developed at NIH to consolidate clinical research data [2]. It is intended to simplify data access to and analysis of data from active clinical trials and to facilitate reuse of existing data to answer new questions. The EHR4CR [3] project supports the feasibility, exploration, design and execution of clinical studies and long-term surveillance of patient populations by providing services supported by an european infrastructure connected to hospital Clinical Data Warehouses. TRANSFoRm project has similar objectives but focused on primary care patient data. STRIDE [4] is a platform supporting clinical and translational research consisting of a clinical data warehouse, an application development framework for building research data management applications and a biospecimen data management system. The ObTiMA system relies on OWL and SWRL to perform semantic mediation between heterogeneous data sources [5]. MATE [6] is an interactive computer system to facilitate explicit, evidence-based decision-making in MDM for

breast cancer care. MATE [7] provides prognostication and risk assessment and also flags up patients eligible for recruiting into ongoing research trials. Trinzeck et al have designed and implemented a generic architecture for Patient Recruitment System compatible with most of the currently available German Hospital Information System environments.

Formal representation of eligibility criteria: The question of the formal representation of eligibility criteria is still an open issue. Several works have been developed or reused formalism for eligibility criteria representation such as CDISC's ASPIRE, Arden Syntax, SAGE, or GELLO or ERGO.

Interoperability and communication: To determine the eligibility status of patient, CRSS has to match clinical data coming from one or several EHR sources to the Clinical trial criteria. This supposes a semantic interoperability and a communication layer between the clinical data sources and the CRSS. To address this issue, the current trend is to combine interoperability standards coming from clinical or research domains (e.g HL7/CDISC), such as the Bridge model.

Methods

Use case addressed by the system: The different steps (S_x) of the MDM process and the system functionalities used all along this process (fig. 1) are : **S1.** A clinical research associate or a principal investigator enter the eligibility criteria of all open clinical trial into the system. This step consist of translating the free text criteria into a computable representation with a graphic user interface. **S2.** Mr. Dupont, a patient having a prostate cancer, consults with Dr. Durant, a urologist. **S3.** Dr Durant schedules a MDM to decide the best therapeutic strategy for his patient. Dr. Durant fills up a MDM form provided by a regional EHR. This MDM report contains the most important medical information about Mr. Dupont, that the MD Team will use during the meeting to take the decision. **S4.** Before the meeting, all the MDM reports are sent through a web service to the system. The system analyses, for each MDM report, all the criteria of all open clinical Trials. **S6.** As output, the system provides an eligibility report for each patient, containing all proposed and rejected clinical trials with the reasoning of the system decision. **S7.** During the meeting, both MDM report and eligibility report are displayed to the multidisciplinary team. The case of Mr. Durant is discussed and a decision is taken by the team. The decision is added to the MDM report. **S8 :** Either, the therapeutic decision is a standard treatment. The MDM report with the decision is then sent to Dr. Durant who informs Mr. Dupont. The patient then either accepts or rejects the proposal. **S9 :** Or according the conclusion of the eligibility criteria report provided by the system, the multidisciplinary team can propose to Mr. Durant to be enrolled into a clinical trial. **S10 :** In this case, a Clinical Research Associate (CRA) is warned after the MDM by a notification. The CRA is in charge to verify details the eligibility of Mr. Durant to. Eventually, Mr. Dupont is contacted, to sign the CT consent. Most of the time, specific exams are scheduled (**S11**) to determine whether the patient is eligible or not. Mr. Dupont is finally included (**S12**).

System Architecture

Patient data source and communication components:

The system was tested with the oncology EHR of the regional oncology network of Brittany. The EHR communicates to the system through secured webs services. A semantic interoperability framework was defined based on the HL7/CDAr2 Level 3 standards and more specifically on HL7 Care Provision/R-MIM Care Record (DSTU) and HL7v3 Standard Transport Specification, Web services Profile, Simple Object Access Protocol version 1.2 (SOAP 1.2). HL7 R-MIM was constrained and only

useful classes for the project were implemented. From this refinement, we have produced a set of HL7 templates to formalise and encode the set of data elements of the MDM reports. To encode both the patient data and the CT criteria, the National Cancer Terminology Thesaurus (NCIT) was chosen as reference ontology/terminology. For each data elements of the MDM report, entities and related value were manually mapped to NCIT concepts. In case of missing concepts, we have enriched the ontology by selecting codes and labels from other reference terminologies (e.g. SNOMED CT, LOINC, etc). The communication between the patient data source and the system was performed by a securitized web service. At the message reception, a virtual Medical Record (vMR) was populated by the system with the transmitted patient data (vMR is a generic HL7 data model dedicated to decision support systems).

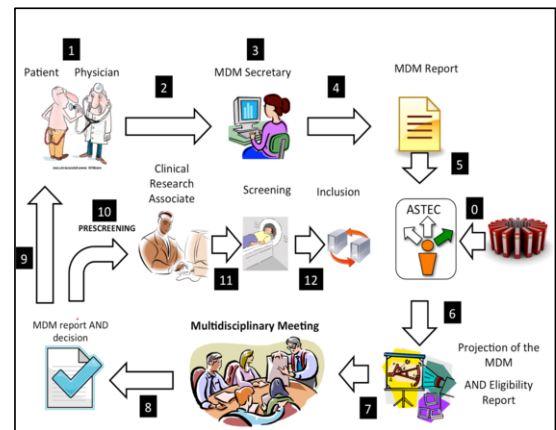


Figure 1 – MultiDisciplinary Meeting workflow

Eligibility and patient data representation: EC refer to complex conditions that a patient needs to fulfill to be included or excluded from a given clinical trial. All clinical trials can be seen as a concatenation of a set of inclusion criteria and a set of exclusion criteria. Patients need to meet the inclusion criteria and avoid the exclusion criteria in order for them to be considered in the clinical trial. Patient data as well as eligibility criteria were formalized relying on OWL models. Namely, patients are defined as a set of *entities*, to which several properties can be associated. Every *property* has a unique *value* that represents its state.

Formalization and processing of the eligible criteria: EC are written by physicians so the terminology used can be subject to interpretation and involve several simple predicates. Thus, a specific criterion can be complex and depend on several *entities*. Complex criteria such as the GETUG 14 eligibility criterion “Cancer in intermediate prognostic group”, were decomposed into a set of *atomic criterion*, i.e. a simple predicate on a unique OWL triplet (e, p, v) and a logic AND/OR relationship. As example, Fig. 2 illustrates how this criterion can be broken down to simple predicates (right side of the graph, i.e., leaves of the graph) and a set of Boolean operations.

A patient is included in a CT if his root criteria score is equal to 1. Only patients that met all inclusion criteria and met none of the exclusion criteria received a score of 1.




Figure 2-...

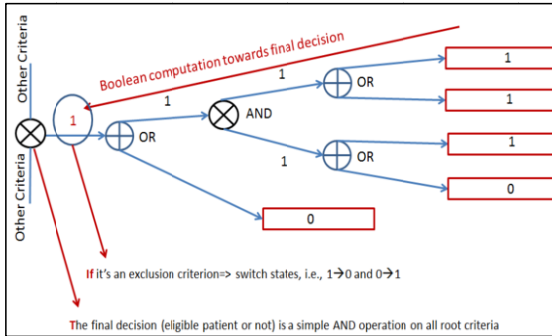


Figure 2- Criterion representation

Inference method: The inference method used in the system has been explained in details in our previous paper [8]. The core of the system is a simple ontology that glues together the components: it imports NCIT ontology and the ontologies that define SWRL built-ins and the classes they could require. This glue ontology defines the classes that are used to represent patient data and the EC in CTs. EC expressed in SWRL are inferred with patient data by an inference engine (JESS). The information granularity is most of the time different between patient data and the EC. This issue is addressed by Jess - subsumption reasoning capability.

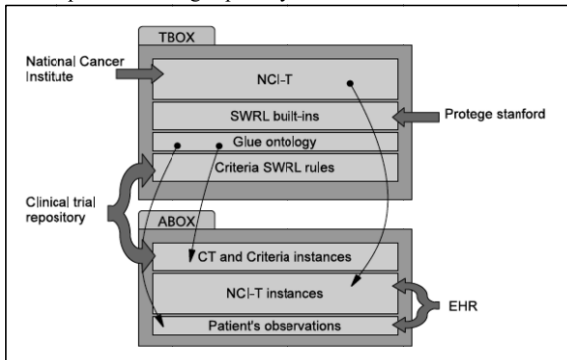


Figure 3 - System design

Figure 3 shows the general architecture: TBox contains the glue ontology that imports NCIT and the SWRL ontologies. The SWRL rules are definitions, and therefore are part of the TBox. The ABox contains the instances of NCIT classes, which are linked by observations, and the instances of the EC.

The eligibility report: A clinical trial has a list of arguments in favour and one against enrolment of the patient. The list in favour is filled by inclusion criteria, while the list against by exclusion criteria. Once the criteria rules are run, the instances of the criteria will have the property *supported_by* either filled with one or more observations, or empty. Inclusion criteria with the property filled are supported, and they in turn support the clinical trial. Exclusion criteria with the *supported_by* property

filled are arguments against the enrolling of the patient to the clinical trial. At this point, rules specific to CTs are run to aggregate the results of the criteria. All inclusion criteria of a CT need to be verified: there must be, for each CT, a rule stating that the conjunction of all inclusion criteria must have at least one supporting argument. It is also possible to trace the exclusion criteria to an observation that fits it. However, extracting the CTs that are supported without arguments against it requires to reason with the closed world assumption: according to the open world assumption, the empty list of arguments against enrolment is considered ignorance, and cannot be used to infer that there are no arguments against enrolment. The last step needs to be performed outside the ontology. An external program obtains the list of all clinical trials and selects those that are both supported and have no arguments against.

Evaluation of the system

Clinical trials for the evaluation of ASTEC

To evaluate our proposed system, we performed retrospective chart review of all patients in every MDMs from October 2008 to march 2009. Only a selection of CTs recruiting at the CEM during the study period was chosen for the ASTEC evaluation. The study focused on prostate cancer, which is the most common cancer in urology and the main topic of discussions at the urology MDM and of many clinical trials. All criteria of the selected CTs were considered for the study, but required some interpretations to be coded in the system. Some high level criteria were rather vague (e.g. "life expectancy ≥ 10 "). Moreover such an item is not present explicitly in the multidisciplinary reports (and probably rarely in most of EHRs). In this case, the Clinical Research Assistant (CRA) in charge of the system setting in accordance with the principal investigator, encoded a set of sub-criteria, that could be match with the information present in the MDM report (to continue with life expectancy ≥ 10 , it was encoded in ASTEC as no hepatic failure AND no renal failure ... AND Age < 75). Eventually, the aggregating algorithm described above, was used to recompose the initial eligibility criteria and determine the final status of eligibility according to the availability of the real data of the MDM reports.

Automatic screening evaluation and Statistical analysis:

We compared screening performed by the ASTEC system with screening performed by two kinds of human decisions:

- The "MDM Decisions", i.e. the real-time decisions taken during MDM, which are finally the real screening decisions;
- The "Reviewer Decisions", i.e. the decisions taken after review by the CRA and the Principal Investigator (PI) of each study. For that, the overall MD reports were reviewed, manually and independently from the MDM decisions, to propose eligible patients to the RCT. Standard performance metrics are used to assess information retrieval: true and false positives, true and false negatives, recall, precision, f1-measure and overall error rate.

$$\text{precision} = \frac{tp}{tp + fp} \quad \text{recall} = \frac{tp}{tp + fn}$$

$$f1 - \text{measure} = 2 \cdot \frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}}$$

Results

During the study period, 23 Multidisciplinary meetings took place with about 285 medical reports. Four ongoing CTs (GETUG 14, 15, 16, 17) concerning prostate cancer were recruiting at the CEM during the study period. Only results from the GETUG 14 are presented, which was the CT with the most eligible patients.




Figure 4 - Venn diagram of the ASTEC decisions, the Reviewer decisions and the MDM decisions

We considered the screening reviewer decisions as gold standard, while the MDM decisions reflected the screening performed in real world by physicians.

As shown in Figure 4, the 3 screening methods all excluded 151 patients. The MDM decisions selected 17 patients eligible to GETUG 14, while the ASTEC system selected 132 patients and the reviewer decisions selected 30 patients eligible to GETUG 14. Finally, the automatic patient screening by the ASTEC system collected 11 additional patients from the selection of the MDM team, and only 2 patients eligible to GETUG 14 were missed.

Table 1-. Performance metrics of the ASTEC system

	MDM Decisions (1)	Reviewer Decisions (2)
n total	285	285
n selected patients	17	30
ASTEC classification results with (1) or (2) as reference		
True positives	17	28
False positives	115	104
True negatives	153	151
False negatives	0	2
Precision	13%	21%
Recall	100%	93%
Error Rate	40%	37%
F1-measure	23%	34%

The ASTEC system had a 21% precision rate and 93% recall rate, and an error rate equal to 37% (**Error! Reference source not found.**).

Error! Reference source not found. presents if each GETUG 14 eligible criterion was verified, not verified, or unknown (when data is missing) for all the screened patients. For instance, only the histological confirmation of the prostate cancer was collected in all the MDM reports.

Table 2- Status of the GETUG 14 eligible criteria

Criteria	Verified	Not verified	Unknown
Histologically confirmed prostate cancer	285 (100%)	0	0
Intermediate-risk cancer	57 (20%)	41 (14%)	187 (66%)
No lymph node invasion	20 (7%)	4 (1%)	261 (92%)
No metastatic disease	14 (5%)	5 (2%)	266 (93%)
PSA < 30 ng/mL	261 (92%)	10 (4%)	14 (5%)
No history of invasive cancer	1 (0.4%)	0	284 (99.6%)
Life expectancy ≥ 10 years	89 (31%)	34 (12%)	162 (57%)
No prior pelvic radiotherapy	27 (9.5%)	2 (0.7%)	256 (89.8%)
No radical prostatectomy (and TURP < 3 months)	34 (11.9%)	5 (1.8%)	246 (86.3%)
No prior hormonal therapy and castration	31 (10.9%)	0	254 (89.1%)
Other exclusion criteria (x5)	171 (12%)	379 (26.6%)	875 (61.4%)

Discussion

Semantic interoperability

We developed an interoperability framework using HL7 standards and the ASTEC ontology (based on the NCIT) to encode the data elements. This approach ensures through a service oriented architecture, an effective and secure communication between the oncology EHR and the CRSS. As far as we know, this is the first CRSS project using this kind of interoperability approach. Despite the fact we tested our system on a single commercial EHR, this architecture will help to connect the system to any kind of data sources using the same communication framework.

Reasoning on structured and coded data

A new contribution of ASTEC is that it demonstrates how it is possible to represent eligibility criteria of clinical trials using SWRL on top of a large domain ontology. Some of the criteria were directly translated into SWRL. Others require more thought, especially these which involve temporal reasoning.

In our project, time of processing is not a stringent requirement: the matching of the patients to the available clinical trials is done offline, before the MDM. The slowest step in the overall procedure is loading NCIT ontology. On a dual core machine, with 8Gb of memory, the process takes over 100 seconds and 2Gb of memory. Importing data into the ontology is nearly instantaneous: we load one patient at the time, and only the clinical trials currently active in the hospital are loaded. The next bottleneck is the conversion of the ontology and of the SWRL rules into Jess, which takes on average 10 seconds. The actual running of the engine takes less than a third of a second. Compared to X¹, we use a much smaller ontology (SNOMED CT is over a million classes, while NCIT is only 75000). Loading the background ontology is performed once.

Integration of ASTEC into the business process of MDM

We have designed ASTEC to be seamlessly integrated into the workflow of MDM. However, in this paper, since we have used a retrospective dataset, we don't provide here any evidence of the user acceptance. This is the object of a current study. Regardless, we can give some arguments in favor. As the data processing occurs before the MDM, the eligibility report is available immediately when the case is discussed by the medical MD Team. The system requires very few interactions from the users. During the MDM, the discussions of the experts are extremely short and efficient, so we don't believe that in case of error or misclassification by ASTEC, the users will have the time to modify and rerun the system on line. ASTEC is not intrusive in the decision process, and it behaves as a reminder system, giving for each patient, a short and clear argumentation in favor or against its selection into a CT.

Evaluation of the performance

The results shows that most of the patients eligible to GETUG 14 were selected by the ASTEC system (28/30 patients). Despite numerous false positives (104 patients), the ASTEC system allowed to rightly exclude 151 patients automatically. The 2 patients missed by the ASTEC system, was also missed by the MDM team.

However, we recognise some limitations. Despite considering 4 different CTs at the beginning of the study, only the results for GETUG 14 are presented here. Indeed, 0, 6 and 1 patients were eligible to GETUG 15, 16 and 17 respectively, while 30 eligible patients were eligible for GETUG 14. This led us to solely discuss ASTEC results on a single clinical trial.

Moreover, GETUG is focused on prostate cancer which is one specific domain in oncology.

The question of the scalability of our system to a broader medical domain is still open. But we can assume that the system will be robust at least for cancer domain. Firstly, because we used quite generic interoperability and reasoning methods. Secondly, because the system is based on a recognized and comprehensive terminology in this domain. And finally, and because the process of MDM is quite similar for all cancers.

Dealing with low data quality and missing data

It is worth noting that decomposing patients' conditions into a set of entities is of course challenging as many entities can be only considered at the physician's discretion. We showed in a preliminary study [9] that a CDSS as the ASTEC system must deal with a high level of missing data and unknown information for the majority of the eligibility criteria. Indeed, we demonstrate the existence of implicit information (i.e. either data not mentioned into MDM reports, or other eligibility criteria than those mentioned by the protocol) used by physicians to perform the screening task. For instance, when a given condition is not present (e.g., heart condition or failure), it is usually omitted and not explicitly set as such into medical record. A physician can overcome this lack of information, but it is hardly the case when it comes to automatic systems. Moreover, we found in the preliminary study a low rate of data quality from the oncology EHR. Initially, we had assumed that multidisciplinary reports were forensic documents, containing mandatory, explicit and comprehensive information. Indeed, the critical decision of the multidisciplinary team is supposed to rely on this set of information. Actually, it turned out that multidisciplinary reports were just a short summary of the main information. They did not cover all the aspects of the patient condition, such as psychological or psychiatric conditions. We also noted that some of very important data which were supposed to be coded by the clinician in the dedicated structured fields of the MDM form, were present but in the free text fields of the reports. So, these data were unusable by the ASTEC system. For instance, in some multidisciplinary reports, TNM staging was not filled out in the dedicated field of the report and was mentioned in the free text but in an implicit way (e.g. "malign tumor of the both sides of the prostate without extension" which implicitly, stands for a T2 stage). It was unlikely for ASTEC, the multidisciplinary team or the reviewer to deduct at a glance the cancer stage and account for this information during decision making.

Such factors explain the main reasons for discrepancies between the system decision and the human decisions, especially based on false positive screened patients by the system. It is nevertheless considered as a decent approximation, one that can enable an efficient pre-selection and help physicians focus on a subset of promising patients.

To address the issue of missing data, we have developed and tested in a related work [8], an OWL model of clinical trial eligibility criteria compatible with partially-known information. In this work, we proposed an OWL design pattern for modeling the eligibility criteria based on the open world assumption. Our approach successfully distinguished clinical trials for which the patient is eligible, clinical trials for which we know that the patient is not eligible and clinical trials for which the patient may be eligible provided that further pieces of information (which we can identify) can be obtained. The OWL study was evaluated on the same clinical trial and on the same set of patients. The results were similar to the ones reported here, with 149 patients potentially eligible patients (132 in the present study). The difference of performance lies in a simpler data extraction process for the OWL model, which does not

affect reasoning. Once the data are extracted from the patient's records, the determination of the patient's eligibility requires a reasoning framework capable of handling both (i) the difference of precision between the data and the criteria, and (ii) the pervasive incompleteness of data. Together, these two studies demonstrate the feasibility of such a task.

Conclusion

Clinical trials are required for the evaluation of medical treatments. Their weakness lies in the difficulty of recruiting enough patients in order to make them statistically meaningful. In this paper we have presented an approach based on OWL and SWRL that addresses the problem of recruitment of patients. The evaluation showed that it is possible to represent the great majority of criteria, and that the system detected most of the patients eligible to the trial and eliminate most of the False negative. The false positives were essentially due to missing data coming from the EHR.

Acknowledgments

The French National Agency of Research (ANR) under the TecSan program 2008 funded the ASTEC project (ASTEC n°ANR-08-003). The authors would like to especially thank Sahar Bayat, Elisabeth Le Pris , Jean-Fran ois Laurent, Nicolas Garcelon, Renaud De Crovoisier for their valuable help.

References

- [1] M. Cuggia, P. Besana, and D. Glasspool, "Comparing semi-automatic systems for recruitment of patients to clinical trials," *Int J Med Inform*, vol. 80, no. 6, pp. 371–388, Jun. 2011.
- [2] James J Cimino and Elaine J Ayres. The clinical research data repository of the us national institutes of health. *Studies in health technology and informatics*, 160(Pt2):1299–1303, 2010.
- [3] G. D. Moor, et al "Using Electronic Health Records for Clinical Research: the Case of the EHR4CR Project," *J Biomed Inform*, Oct. 2014.
- [4] H. J. Lowe, T. A. Ferris, P. M. Hernandez, and S. C. Weber, "STRIDE – An Integrated Standards-Based Translational Research Informatics Platform," *AMIA Annu Symp Proc*, vol. 2009, pp. 391–395, 2009.
- [5] H Stenzhorn, et al. The ObTiMA system - ontology-based managing of clinical trials. *Studies in health technology and informatics*, 160(Pt 2):1090–1094, 2010.
- [6] P Vivek, et al "Cancer Multidisciplinary Team Meetings: Evidence, Challenges, and the Role of Clinical Decision Support Technology." *International Journal of Breast Cancer* 2011 (12 2011): 1-7. doi:10.4061/2011/831605.
- [7] B. Trinczek, « Design and multicentric implementation of a generic software architecture for patient recruitment systems re-using existing HIS tools and routine patient data », *Appl Clin Inform*, vol. 5, n° 1, p. 264 283, 2014.
- [8] O. Dameron, et al. , "OWL model of clinical trial eligibility criteria compatible with partially-known information," *J Biomed Semantics*, vol. 4, no. 1, p. 17, 2013.
- [9] B Campillo-Gimenez, et al. Improving the pre-screening of eligible patients in order to increase enrollment in cancer clinical trials. *Trials* 2015- accepted

Author Correspondance

Marc CUGGIA (marc.cuggia@univ-rennes1.fr)