



Published in final edited form as:

J Biomed Inform. 2015 December ; 58(Suppl): S183–S188. doi:10.1016/j.jbi.2015.06.013.

Using Local Lexicalized Rules to Identify Heart Disease Risk Factors in Clinical Notes

George Karystianis^{1,5}, Azad Dehghan¹, Aleksandar Kovacevic², John A. Keane^{1,4}, and Goran Nenadic^{1,3,4}

¹School of Computer Science, University of Manchester, Manchester, UK

²Faculty of Technical Sciences, University of Novi Sad, Novi Sad, Serbia

³Health eResearch Centre, Manchester, UK

⁴Manchester Institute of Biotechnology, University of Manchester, UK

⁵The Christie NHS Foundation Trust, Manchester, UK

Abstract

Heart disease is the leading cause of death globally and a significant part of the human population lives with it. A number of risk factors have been recognised as contributing to the disease, including obesity, coronary artery disease (CAD), hypertension, hyperlipidemia, diabetes, smoking, and family history of premature CAD. This paper describes and evaluates a methodology to extract mentions of such risk factors from diabetic clinical notes, which was a task of the i2b2/UTHealth 2014 Challenge in Natural Language Processing for Clinical Data. The methodology is knowledge-driven and the system implements local lexicalised rules (based on syntactical patterns observed in notes) combined with manually constructed dictionaries that characterize the domain. A part of the task was also to detect the time interval in which the risk factors were present in a patient.

The system was applied to an evaluation set of 514 unseen notes and achieved a micro-average F-score of 88% (with 86% precision and 90% recall). While the identification of CAD family history, medication and some of the related disease factors (e.g. hypertension, diabetes, hyperlipidemia) showed quite good results, the identification of CAD-specific indicators proved to be more challenging (F-score of 74%). Overall, the results are encouraging and suggested that automated text mining methods can be used to process clinical notes to identify risk factors and monitor progression of heart disease on a large-scale, providing necessary data for clinical and epidemiological studies.

Corresponding author g.nenadic@manchester.ac.uk, School of Computer Science, Oxford Road, Manchester M13 9PL, United Kingdom, Tel. +44-161-2756289.

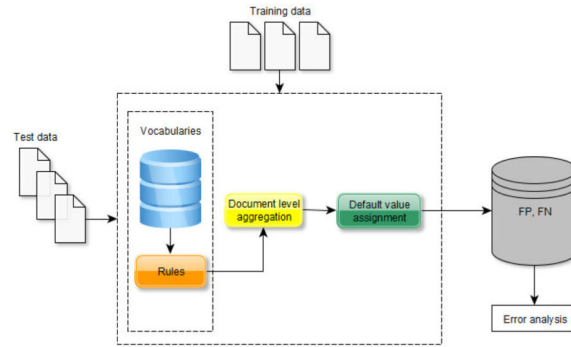
Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Availability

The resources developed as part of the system are available at <http://gnode1.mib.man.ac.uk/tools/i2b2-2014-task2/>.

The authors declare no conflict of interest.

Graphical abstract



Keywords

Text mining; risk factors; heart disease; vocabularies; rule-based modelling

1. Introduction

Heart disease is the leading cause of death globally¹: in the UK, for example, about one in six men and one in ten women die from heart disease². Furthermore, a significant part of the human population lives with it (e.g., 2.3 million people in the UK alone). Many studies have been conducted to improve treatment and identify possible risk factors and life-style habits that may make a person more likely to develop heart disease. For example, obesity, coronary artery disease (CAD), hypertension, hyperlipidemia, diabetes, smoking, family history of premature CAD, unhealthy diet and age above 55 have been acknowledged as important risk factors³. The ability to identify such risk factors for individual patients is important for both disease prevention and treatment; furthermore, extraction of such information on a large-scale (e.g., from electronic health records (EHRs)) is key for epidemiological studies and understanding the development of the disease.

While EHRs contain coded (structured) information that is undoubtedly useful for such studies, clinical narratives (e.g., letters, doctor notes) are in an unstructured, free-text form and often include rich, contextual information that is not present elsewhere. Processing of such information has been a focus of clinical text mining for over 30 years⁴⁻⁶, with notable results in harvesting important clinical concepts and events. Efforts have focused on the identification of various concepts, combining a variety of approaches. For example, Goryachev et al⁸ recognised family history from discharge summaries and outpatient clinic notes through a rule-based approach with an F-score of 95%, while Wang⁹ extracted findings and medical procedures in clinical progress notes by applying both a rule-based system and modelling a conditional random field classifier, with F-scores of 49% and 82% respectively. Several approaches have been developed for identification of medication information from clinical notes. Patrick et al¹⁰, for example, applied a hybrid approach of supervised learning and rules, while Spasic et al¹¹ used a rule-based methodology and Yang¹² mainly relied on a dictionary-based method. Other work has focused on the extraction of medical problems, treatments and tests from clinical narratives with relatively

good results, typically with an F-score of around 80%. For example, Rink et al¹³ and Jonnalagadda et al¹⁴ applied machine learning, whereas Xu et al¹⁵ combined machine learning and rules for that task. Finally, there has been previous work on extracting risk factors for certain conditions: for example, Fiszman and colleagues¹⁶ used a semantic processor that recognised predications to extract metabolic syndrome risk factors (such as obesity, high density lipoprotein, elevated blood pressure) from MEDLINE abstracts, with an overall F-score of 59%.

Several community challenges in clinical text processing have been organised to assess the state-of-the-art for specific tasks, including, for example, medication identification¹⁷, extraction of co-morbidities¹⁸, etc. One of the tasks in the 2014 i2b2/UTHealth Challenge in Natural Language Processing for Clinical Data aimed to identify potential risk factors for heart disease from clinical notes of diabetic patients³. The task focused on eight classes, including mentions of *CAD* or factors that are associated with its onset (*diabetes*, *obesity*, *hyperlipidemia*, *hypertension*, *smoking status*, *family history* of *CAD* and related *medications*). In this paper we describe and evaluate our approach to the task, which uses local lexicalised rules combined with manually constructed dictionaries that characterize the domain. We demonstrate that the rule-based approach is feasible and can be used for reliable large-scale data harvesting.

2. Materials and Methods

Task and Data

The task focused on document-level extraction of the eight classes listed above, where each class is characterized by attributes (see Table 1). The five disease classes (*CAD*, *diabetes*, *obesity*, *hyperlipidemia*, *hypertension*) are recognised through either explicit presence (mention) of the disease or the progression of clinical markers suggesting the targeted disease (e.g., "hemoglobin levels above 6.5" and "glucose levels over 126" are indicators for *diabetes*). Different diseases have a different number of indicators: for example, *CAD* has four (a *mention* of the disease or its *symptom* - e.g., *angina*, *event* - e.g., *heart catheterization*, or *test* - e.g., *stress test*), *obesity* has three (*mention*, *body mass index (BMI)* and *waist circumference (WC)*), while *hypertension* has two (*mention* and *high blood pressure*).

The medication class has two attributes: *type1* is the drug category to which the medication belongs (a total of 22, e.g., "sulfonylureas", "meglitinides") and *type2* indicates drugs that can be included in more than one category (e.g., "zestoretic" has *type1* of "ACE inhibitor" and *type2* of "diuretic").

The time attribute refers to the temporal interval in which a risk factor was present in the patient's medical history: *before* the Document Creation Time (DCT, i.e. the time when the clinical note was created), *during* DCT and/or *after* DCT. DCT is considered as an attribute in all of the disease factors and in the medication class. We note that a specific risk factor can be present before, during and after DCT, or in any combination of these.

The smoker class has a *status* attribute that indicates whether the person is a "current", "past", "ever", or "never" smoker, or if their smoking status is "unknown". Finally, the family history contains the "present" or "not present" indicator that specifies whether the patient has first degree relatives (e.g., parents, siblings) who were diagnosed prematurely with CAD.

The overall task was to indicate the presence of these risk factors the document level. Specifically, for the five disease factors, the task included a binary, document-level classification (present/absent) for each of the associated indicators and also for the explicit disease mentions. The time attribute further specifies the timeframe(s) (before, during, after). Medication information includes the two types and time (3 values), whereas family history of CAD is a binary classification task (present/absent). Finally, the smoking status needs to be instantiated with one of the five possible values. The task organizers provided a training set (790 clinical notes) and 514 notes as an evaluation set, all fully annotated at the document level³. The data are available at the following link <https://www.i2b2.org/NLP/DataSets>.

Method overview

After an initial analysis of the training set where we observed common lexical patterns that indicate the presence of the targeted factors (e.g., "male with hypertension", "pmh: diabetes, hypertension"), we designed and implemented a lexicalized rule-based approach for their recognition. Our methodology consists of four steps:

- **Step 1:** creation of specific vocabularies for each class.
- **Step 2:** design and implementation of rules to capture risk factors of interest at the mention level.
- **Step 3:** integration of the mention-level results at the document level.
- **Step 4:** designating the time value to the identified factors.

In the first step, a number of task-specific semantic groups have been identified and lexicalized through a set of custom-made vocabularies that were engineered from open clinical resources (see Table 2). The dictionaries were manually tailored by observing the training set for the usage of terms describing the associated risk factors and expressions related to their indicators (e.g., "blood pressure", "high blood pressure", "systolic blood pressure", etc.), and by adding clinical synonyms and acronyms from the Unified Medical Language System²¹ (UMLS) for specific terms of interest.

In the second step, these dictionaries are used to anchor and constrain a generic set of local rules for identification of disease and risk factor mentions by using:

- specific semantic groups, recognised by the vocabularies and/or regular expressions,
- semi-frozen lexical expressions (e.g. "*the patient was diagnosed with*") that are used as anchors for specific entity and attribute types.

We note that the vocabularies used are task-specific, whereas the general rule patterns are focused on the identification of disease risk factors, which are then used to "infer" the

mention of a specific disease type (e.g., based on the specific indicators). Generally speaking, the rules were based on structured patterns consisting of semi-frozen syntactic chunks (e.g. noun phrases, verbs and prepositions) and/or semantic place holders (through dictionary mentions), either suggesting the presence of a disease or associated event (e.g., “history of present illnesses include hypertension”, “underwent catheterization”, “stress test was positive”) or specific indicators (e.g., via specific measurements implemented through regular expressions (e.g., “BMI: NUMBER”, “blood pressure: NUMBER/NUMBER”). For example, a rule that captures mentions of CAD-specific surgery would have two parts: a semi-frozen verbal expression (e.g. various forms of “undergo”) followed by a mention matched by the surgery dictionary (as mentioned in Table 2). We have also implemented concept enumeration as it appears quite frequently in the training data, particularly for disease and medication mentions (e.g., “pmhx: *dm, htn, dementia*”, “Medications: *Lisinopril Pravachol*”). Table 3 presents examples of rules for some risk factors. The number of rules for specific entity types (see Table 4) roughly indicates the complexity of the targeted information i.e., the number of associated indicator types. For the design and implementation of the rules we used MinorThird²², an information extraction development environment that we have previously used for clinical text mining¹¹.

We note that a document in this task was a set of clinical notes for a given patient and that we are interested whether a risk factor is mentioned or not within the document. Therefore, in the third step, we have integrated the data identified at the mention level to the document level. For example, if we have detected any high blood pressure indicators (e.g., “*bp 150/90 mm/hg*” or “*blood pressure: 160/90 mm/hg*”) in a note, we consider that the patient has “hypertension”, with an indicator of “high blood pressure” tagged at the document level.

This approach was followed for all entity types and attributes apart from the time dimension. As the clinical notes were longitudinal, there was a high chance that patients have a number of diseases (and indicators) before, during and after the DCT. This is also likely to be the case with the (majority of) administered medications. This hypothesis was confirmed by the training set: from 1,223 disease mentions (at the document level), only 15 (1.23%) did not have all three time attributes values (i.e. before DCT, during DCT, after DCT); for medication mentions, from a total of 2,191, only 203 (9.26%) had either one or two time attributes. Therefore, we decided to set, as a default, all three values for the time attribute for all the disease and medication mentions, and aim to identify only explicit localized expressions (e.g., “stop *drug*”; “start on *drug*”) to alter these if necessary. Specific disease indicators were treated by different defaults. For example, *body mass index* and *high blood pressure* are typically recorded during the creation of the narrative and rarely denote past or future values; we therefore decided to assign the default value of “during DCT” to their time attribute. Other tests (e.g. levels of hemoglobin, glucose, high LDL, high cholesterol, CAD test, CAD symptom, CAD event) typically happen before the date of the current note and hence were set to the default value of “before DCT”. This was also supported by the data in the training corpus.

3. Results

The system was formally evaluated as part of the i2b2 challenge. Table 5 displays the summarized results across all data sets (training, development and evaluation). The overall micro precision was 85.57% with recall of 90.07% and a micro F-score of 87.76%. We note that there was only a marginal drop in the performance compared to the training data, suggesting that the lexicalised rules managed to generalise the risk factor identification quite well. Table 6 shows the results per entity class for the evaluation set. The highest F-score was returned for family history (95.91%), with the highest recall of 96.97% for medication. With the exception of CAD which proved to be the most challenging class to recognize (F-score of 73.63%), all other classes were identified with an F-score above 85% indicating that the approach we followed was effective in the identification of several components of CAD risk factors in clinical narratives.

4. Discussion

The system's micro F-score of 87.76% ranked the system 9th out of the 48 submissions (up to 3 submissions per a team). We note that the performance of the rule-based approach was well above the challenge mean (81.5%) and 5% less than the highest ranking system. This suggests that a rule-based approach for the recognition of heart disease risk factors and the assignment of a time indicator regarding their progression (or not) is worthwhile. To perform an analysis of false positives (FPs) and negatives (FNs), we took a random sample of ten documents for each class from the evaluation set and observed the common types of error that the system generated.

False positives

A quarter (10/48) of FPs originated from disease mentions that are either related to the family of the patient (e.g., "*family history: diabetes, fh: -dm: father*"), are negated (e.g., "*no history of hypertension*") or refer to allergies (e.g. "*Allergies: sulfa drugs*"). It is interesting that negation was not that frequent, contributing to only 6% of cases. Another quarter of FPs (11/48) were ambiguous cues: for example, "*lipids*" and "*ht*" are often used to describe the diagnosis and the status of hyperlipidemia and hypertension respectively, but they can also be used in a different context (e.g., "*lipids will be checked*", "*ht 1.82 cm*"). "*Insulin*" was another frequent example, as it refers to both disease mentions (e.g., "*insulin dependent diabetes mellitus*") and a medication. We note that over half of FPs (27/48) are risk factors possibly missed in the annotation process (e.g. "*hemoglobin a1c 7.7*" was not annotated as an indicator for CAD; similarly "*abd: obese, non-protuberant*", "*hbalc 09/20/2065 6.50*", "*glu 200-265*", "*obese older gentleman*", "*medications: baclofen, atenolol, lactulose and lasix*").

False negatives

In a number of cases, the system ignored disease mentions in particular. For example, CAD attributes (*event, test, mention, and symptom*) were much more variable compared to other classes (e.g., hypertension or hyperlipidemia) and a number of mentions were missed as the rules (although the largest in number) were not flexible enough (25 out of 65 cases) or

lexical/variation coverage was limited (e.g. unknown abbreviations, 16/65 cases: for example, "*3 vessel coronary artery bypass surgery*" has appeared in a number of variants, including "*3-vessel ca bypass surgery*", "*3-v ca bypass surgery*" or "*3v bypass surgery*"). This suggests that an extension of the vocabularies could lead to an improvement towards the system's performance. Furthermore, some of the rules used enumerations of diseases (e.g., "*medical issues include list-of-diseases*"), but the cases where a mention was not recognised (e.g., "*diverticulosis*", "*seronegative ra*") would trigger a termination of the enumerated list (and thus the end of the rule match) and as a result a number of mentions were missed. Finally, some particular indicators required clinical background knowledge (e.g., "**LDL 111**" as a high LDL indicator for hyperlipidemia), which was not encoded.

Time attribute

The implementation of the time attribute default values for the medication and disease mentions has also contributed to some FPs and FNs. As expected, due to the application of the default rule (assigning all three time attribute values to disease and medication mentions), we found FP time attribute values for disease mentions only in nine out of 514 notes (1.75%). In addition, for medication mentions, we detected 144 documents containing FPs, resulting in a lower precision for medications (82%) when compared to the other classes. Although we have implemented some exceptions (e.g., "*stop drug*", "*start on drug*"), there were cases where these further required handling medication enumerations (e.g., in "*Patient was immediately told to stop both her Roxicet and Monopril*" we correctly time-framed *Roxicet* but *Monopril* was on the default rule, making both [Monopril, during DCT] and [Monopril, after DCT] false positives). Still, this approach contributed to the highest recall (97%) for the medication mentions. Overall, the decision to implement default rules for the time attribute appears to be justified. Although the number of errors generated was not large, more sophisticated temporal information²³ could have contributed to the increase of the system's performance.

While the design and implementation of rule-based systems is known to be time consuming, in this case the whole system was engineered within ~6 weeks FTE (full-time equivalent), with the system fully operational within a month with further tests aiming to improve its efficiency in the remaining time. We have combined different expertise within the team, covering both clinical aspects and text mining experience, which allowed for a rapid domain-driven development of lexicalized rules. We purposely separated designing the common syntactical patterns for the identification of risk factors from the lexical modelling; therefore, the system can be tailored for the recognition of other targeted mentions by providing the necessary vocabularies, possibly from the existing clinical terminologies. Nonetheless, a significant number of rules involved the identification and "interpretation" of specific (measured) indicators (e.g. "LDL 111" is an indicator of high cholesterol); such rules will require redevelopment in case of a new task and potential linking to a clinical knowledge base.

5. Conclusions

The objective of the i2b2 2014 task was to recognise heart disease risk factors from clinical narratives of diabetic patients and assign the respective time intervals. In this paper we have described a methodology that is based on local rules lexicalized with extensive vocabularies that represent specific classes. The mention-level results were aggregated at the document level. The time attribute for each class relied on a number of specific default values. The overall performance of 88% F-score suggests that a lexicalised rule-based approach combined with default values can be used to process clinical notes to identify risk factors and monitor progression of heart disease on a large-scale, providing necessary data for clinical and epidemiological studies.

Future work includes the implementation of temporal extraction that will assist in the assignment of time values. Identification of a wider context of a disease or medication mention (e.g. relevant section (history, directions, course of treatment, allergies) and whether the mention refers to an event that is questionable/planned/negated are other areas that can contribute to better system performance. Finally, adding a clinical knowledge base and, in a real-world settings, the use of structured data (e.g. test/laboratory results) that appear in EHRs is a potential approach that can be used for data integration, validation and consolidation.

Acknowledgments

This work has been supported by Health e-Research Centre (HeRC), The Christie Hospital NHS Foundation Trust, the KidsCan charity, the Royal Manchester Children's Hospital and the Serbian Ministry of Education and Science (projects III44006; III47003).

References

1. World Health Organization. The top 10 causes of death. <http://www.who.int/mediacentre/factsheets/fs310/en/>
2. Shah ASV, Griffiths M, Ken KL, McAllister DA, Hunter AL, Ferry AV, Cruikshank A, et al. High sensitivity cardiac troponin and the under-diagnosis of myocardial infarction in women: prospective cohort study. *BMJ*. 2015; 350:g7873. [PubMed: 25609052]
3. Stubbs A, Kotfila C, Xu H, Uzuner O. Practical applications for NLP in Clinical Research: the 2014 i2b2/UTHealth shared tasks. 2015 (Submitted).
4. Friedman C, Shagina L, Lussier YA, Hripcsak G. Automated encoding of clinical documents based on natural language processing. *J Am Med Inform Assoc*. 2004 Sep-Oct;11(5):392–402. Epub 2004 Jun 7. [PubMed: 15187068]
5. Savova GK, Masanz JJ, Ogren PV, Zheng J, Sohn S, Kiper-Schuler KC, Chute CG. Mayo clinical Text Analysis and Knowledge Extraction System (cTAKES): architecture, component evaluation and applications. *J Am Med Inform Assoc*. 2010; 17.5:507–513. [PubMed: 20819853]
6. Spasi I, Livsey J, Keane JA, Nenadic G. Text mining of cancer-related information: review of current status and future directions. *Int J Med Inform*. 2014 Sep; 83(9):605–623. [PubMed: 25008281]
7. Sohn S, Kocher JPA, Chute CG, Savova GK. Drug side effect extraction from clinical narratives of psychiatry and psychology patients. *J Am Med Inform Assoc*. 2011 Dec; 18(Suppl 1):i144–i149. [PubMed: 21946242]
8. Goryachev, S.; Hyeoneui, K.; Qing, ZT. AMIA Annual Symposium Proceedings. Vol. 2008. American Medical Informatics Association; 2008. Identification and extraction of family history information from clinical reports.

9. Wang, Y. Proceedings of the ACL-IJCNLP 2009 Student Research Workshop. Association for Computational Linguistics; 2009. Annotating and recognising named entities in clinical notes.
10. Patrick J, Li Min. High accuracy information extraction of medication information from clinical notes: 2009 i2b2 medication extraction challenge. *J Am Med Inform Assoc.* 2010; 17:524–527. [PubMed: 20819856]
11. Spasic I, Sarafraz F, Keane AJ, Nenadic G. Medication information extraction with linguistic pattern matching and semantic rules. *J Am Med Inform Assoc.* 2010; 17:532–535. [PubMed: 20819858]
12. Yang H. Automatic extraction of medication information from medical discharge summaries. *J Am Med Inform Assoc.* 2010; 17:545–548. [PubMed: 20819861]
13. Rink B, Harabagiu S, Kirk R. Automatic extraction of relations between medical concepts in clinical texts. *J Am Med Inform Assoc.* 2011; 18.5:594–600. [PubMed: 21846787]
14. Jonnalagadda S, Cohen T, Wu S, Gonzalez G. Enhancing clinical concept extraction with distributional semantics. *Journal of Biomedical Informatics.* 2012; 45.1:129–140. [PubMed: 22085698]
15. Xu Y, Hong K, Tsujii J, Chang I-Chao E. Feature engineering combined with machine learning and rule-based methods for structured information extraction from narrative clinical discharge summaries. *Journal of the American Medical Informatics Association.* 2012; 19:824–832. [PubMed: 22586067]
16. Fiszman, M.; Rosemblat, G.; Ahlers, CB.; Rindfleisch, TC. AMIA Annual Symposium Proceedings. Vol. 2007. American Medical Informatics Association; 2007. Identifying risk factors for metabolic syndrome in biomedical text.
17. Uzuner O, Solti I, Cadag E. Extracting medication information from clinical text. *J Am Med Inform Assoc.* 2010; 17.5:514–518. [PubMed: 20819854]
18. Uzuner Ö. Recognizing Obesity and Comorbidities in Sparse Data. *J Am Med Inform Assoc.* 2009; 16(4):561–570. [PubMed: 19390096]
19. Deleger L, Grouin C, Zweigenbaum P. Extracting medication information from narrative patient records: the case of medication-related information. *J Am Med Inform Assoc.* 2010 Sep-Oct;17(5): 555–558. [PubMed: 20819863]
20. Doan S, Colier N, Xu H, Duy HP, Phuong MT. Recognition of medication information from discharge summaries using ensembles of classifiers. *BMC Med Inform Decis Mak.* 2012; 12:36. [PubMed: 22564405]
21. UMLS. 2014. <http://www.nlm.nih.gov/research/umls/>
22. Cohen, WW. MinorThird: Methods for identifying names and ontological relations in text using heuristics for inducing regularities from data. 2004. <http://github.com/TeamCohen/MinorThird/>
23. Kovacevic A, Dehghan A, Filannino M, Keane J, Nenadic G. Combining rules and machine learning for extraction of temporal expressions and events from clinical narratives. *J Am Med Inform Assoc.*

Highlights

- We created a set of task-specific dictionaries related to heart disease.
- We designed generic rules for risk factors identification.
- The result are aggregated at the document level
- Temporal attributes are assigned class-specific defaults.
- Rule-based risk factor extraction is feasible and reliable.

Table 1

Heart disease risk factors with their attributes as used in the challenge. The table includes examples and the number of mentions in the training set. Indicators are specific markers that indicate the factor's presence in a patient. Time suggests the period in which the risk factor was present with regards to the creation of the clinical note. The bold parts indicate the targeted mentions.

Risk factor	Attributes		Example	Number of mentions
	Indicator	Time		
hyperlipidemia	disease mention	before, during, after DCT	“PMH: S/p mechanical aortic valve replacement, CHF, HTN, hyperlipidemia ”	340
	high cholesterol		“patient's Chol 179 ”	5
	low-density lipoprotein		“ LDL 119 ”	33
hypertension	disease mention	before, during, after DCT	“Medications include for hypertension , diabetes, and hypercholesterolemia”	524
	high blood pressure		“ Blood pressure: 150/92 ”	33
diabetes	disease mention	before, during, after DCT	“PAST MEDICAL HISTORY: Remarkable for seizure, type II diabetes mellitus , panhypopituitarism secondary”	524
	glucose levels		“ glu 192 ”	24
	haemoglobin levels		“ Hemoglobin a1c 8.3 ”	101
obesity	disease mention	before, during, after DCT	“Diabetes mellitus Hypertension Obesity ”	147
	bmi		“ BMI 30.3 ”	16
	waist circumference		-	0
CAD	disease mention	before, during, after DCT	“In addition, CAD , diabetes, hypertension, CHF”	261
	event		“She was treated for NSTEMI with ASA 325 mg”	237
	test		“and a stress test suggesting anterior ischemia”	74
	symptom		“cardiac catheterization laboratory because of progressive worsening angina ”	68
medication	type1	before, during, after DCT	“Medications Lisinopril ”	3,085
	type2		“Medications Avandamet ”	13
family history of CAD	present, not present		“ Mother diagnosed with cad ”	22
smoker	status – current		Currently smoking a pack per day	57
	status – past		Ex-smoker	149
	status – never		“ smoking - no ”	184
	status - ever		He smoked once	9
	status – unknown		-	373

Table 2

Dictionaries used for the lexicalisation of rules. A total of 21 dictionaries were manually curated

Dictionary	Example terms	size
haemoglobin	hgblc, hemoglobin, glycohemoglobin, hbg	14
diabetes	type 2 diabetes, insulin dependant diabetes, non-insulin-dependant diabetes, adult onset diabetes	66
hyperlipidemia	hld, hypercholesterolemia, hyperlipoproteinemia, dyslipidemia	14
CAD	cad, coronary artery disease, three-vessel coronary artery disease, heart disease	11
hypertension	htn, essential hypertension, hypertension, hypertensive disorder	9
CAD symptom	chest pressure, angina, substernal chest pain, intermitten angina, mild chest discomfort	40
myocardial infarction	anteroseptal mi, lateral myocardial infarction, prior inferior myocardial infarction	49
surgery	angioplasty, coronary artery bypass, cardiac bypass graft surgery, poba, 2v cabg	60
smoking concepts	tobacco use, cigarette smoking, tobacco smoking, cigarette abuse, cigar abuse	28
former smoker	former smoker, ex smoker, prior smoker, remote smoker, former heavy smoker	10
obesity	central obesity, adiposity, obese, general obesity, obesity, morbid obesity,	7
blood pressure	blood pressure, bp, sbp, dbp, blood pressures, systolic blood pressure, hbp	7
gender	lady, gentleman, man, woman, patient, male, female, f, m	9
history	past medical history, pmh, pmhx, history, background history, previous history	10
social activity	alcohol consumption, alcohol use, substances, substance abuse, drinking, narcotics	17
medication head	ointment, inhaler, nebulizer, nebs, puffer, sulphate, cream, paste, elixir, lotion	47
CAD relative	brother, mother, father, sister, children, son, daughter,	7
catheter	left heart catherization, cardiac catherization, cardiac cath	10
CAD stent	RCA stent, aterial stent, taxus stent, cardiac stent, right coronary stent, cypher stent	10
CAD test	stress test, stress mibi, thallium stress, exercise tolerance test, mibi	5
diseases	osteoarthritis, depression, Parkinson's disease, glaucoma, attention deficit disorder	100

Table 3

Examples of rules for the recognition of heart disease risk factors.

Examples show both an “abstract” description of the rule and the MinorThird notation. Rule components in square brackets are the extracted (target) spans that denote the mention of interest; the rest of the rule (if any) specifies the context. The rules use explicit matching of spans (e.g., eq('past') matches string 'past'), regular expressions (re) for matching specific frozen expressions and clues and the vocabularies that contain mentions of specific dictionaries. For example, @surgery includes various surgical procedures, @pressure has variations of blood pressure and @history contains expressions that suggest a mention of history of disease (see Table 2). "Any" matches a given number of tokens (e.g. any{0,4} matches up to 4 tokens).

class	Examples	identified span			
		past-medical-history (NP)	any token	preposition	disease mention (NP)
diabetes (disease mention)	abstract rule	@history	any{0,4}	re('(offor)')	[@diabetes]
	rule example	past medical history	is also positive	for	non-insulin diabetes mellitus
CAD (event)	abstract rule	undergo (verb)	any token	CAD related surgery (NP)	
	example rule	re('undergone underwent undergo')	any{0,1}	[@surgery]	
Hypertension (blood pressure)	abstract rule	underwent	a	3-vessel coronary artery bypass surgery	
	example rule	Since I saw Ms. Law, she underwent a 3-vessel coronary artery bypass surgery			
Hypertension (blood pressure)	abstract rule	blood pressure (NP)	numeric regex	punctuation	numeric regex
	example rule	[@pressure]	re('[1-3]?[0-9][0-9]')	a(punctuation)	re('[1-3]?[0-9][0-9]')
Hypertension (blood pressure)	abstract rule	blood pressure	192	/	94
	example rule	Vital signs: blood pressure 192/94			
Hyperlipidemia (cholesterol)	abstract rule	gender (NP)	past-medical-history (NP)	disease mention (NP)	disease mention (NP)
	example rule	@gender	@history?	@disease?	[@hyperlipidemia]
Hyperlipidemia (cholesterol)	abstract rule	man	with	cad, dm,	high cholesterol
	example rule	He is a 61-year-old man with cad, dm, high cholesterol, hm and family history of early cad.			

Table 4

The number of rules created for each of the targeted risk factors.

Risk factors	Number of rules
medication	10
hyperlipidemia	66
hypertension	70
diabetes	91
obese	63
CAD	133
family history of CAD	21
smoker	94

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 5

Results per data sets.

The training data (790 notes) were distributed in two batches (an initial set of 521 notes, followed by a development set of 269 notes). While the initial training dataset was used for rule engineering and building of lexical resources, the development set was used for internal validation during the implementation.

Data	Micro		
	precision	recall	F-score
Initial training set (521 notes)	85.64	92.63	89.01
Development set (269 notes)	83.88	91.84	87.68
Evaluation set (514 notes)	85.57	90.07	87.76

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 6

Results per risk factor class in the evaluation set.

class	frequency	Micro		
		precision	recall	F-score
Obesity	262	83.15	86.64	84.66
Diabetes	1,189	93.27	79.83	86.03
CAD	1,021	78.11	69.64	73.63
Hypertension	1,308	95.53	85.92	90.47
Hyperlipidemia	751	90.94	82.82	86.69
Family history	514	95.91	95.91	95.91
Smoker	514	85.21	85.55	85.38
Medication	5,825	82.24	96.97	89.00
Overall (run 1)	11,384	85.57	90.07	87.76

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript