Open Research Online

The Open University's repository of research publications and other research outputs

Reasoning on health condition evolution for enhanced detection of vulnerable people in emergency settings

Conference or Workshop Item

How to cite:

Morales Tirado, Alba Catalina; Daga, Enrico and Motta, Enrico (2021). Reasoning on health condition evolution for enhanced detection of vulnerable people in emergency settings. In: K-CAP '21: Proceedings of the 11th on Knowledge Capture Conference, 2-3 Dec 2021, Virtual Event, USA, ACM, pp. 9–16.

For guidance on citations see [FAQs](#).

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

<https://creativecommons.org/licenses/by-nc-nd/4.0/>

Version: Accepted Manuscript

Link(s) to article on publisher's website:
<http://dx.doi.org/doi:10.1145/3460210.3493551>

Copyright and Moral Rights for the articles on this site are retained by the individual authors and/or other copyright owners. For more information on Open Research Online's data [policy](#) on reuse of materials please consult the policies page.

oro.open.ac.uk

Reasoning on Health Condition Evolution for Enhanced Detection of Vulnerable People in Emergency Settings

Alba Catalina Morales Tirado
KMi, The Open University
Milton Keynes, United Kingdom
alba.morales-tirado@open.ac.uk

Enrico Daga
KMi, The Open University
Milton Keynes, United Kingdom
enrico.daga@open.ac.uk

Enrico Motta
KMi, The Open University
Milton Keynes, United Kingdom
enrico.motta@open.ac.uk

ABSTRACT

During an emergency event, such as a fire evacuation, support services benefit from having information about people who may require special assistance. In this context, health data represents a particularly important source of information, as it can allow an emergency response system to build an accurate picture of people's relevant health conditions and use this to advise responders. However, to perform this task, a system needs to represent and reason over the evolution of health conditions over time. Crucially, it needs to predict the probability that a potentially relevant condition mentioned in a health record is still valid at the time of the emergency. In this paper, we propose a methodology for representing the evolution of health conditions and reasoning about them in the context of an emergency scenario. To support our approach with data, we develop a pipeline to capture knowledge about condition evolution from reliable sources in natural language. We incorporate these two components into a system that predicts a person's likelihood of being vulnerable during an emergency event. Finally, we demonstrate that representing and reasoning about condition evolution improves the quality and precision of the recommendations provided by our system to emergency services.

CCS CONCEPTS

• **Computing methodologies** → **Knowledge representation and reasoning**; **Information extraction**; • **Information systems** → *Expert systems*.

KEYWORDS

Health Records, Emergency Services, Datasets, Text Classification, SNOMED CT

ACM Reference Format:

Alba Catalina Morales Tirado, Enrico Daga, and Enrico Motta. 2021. Reasoning on Health Condition Evolution for Enhanced Detection of Vulnerable People in Emergency Settings. In *Proceedings of the 11th Knowledge Capture Conference (K-CAP '21)*, December 2–3, 2021, Virtual Event, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3460210.3493551>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
K-CAP '21, December 2–3, 2021, Virtual Event, USA.

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-8457-5/21/12...\$15.00
<https://doi.org/10.1145/3460210.3493551>

1 INTRODUCTION

Healthcare data provides a significant and valuable source of information and, in particular, its use has received increased attention in the context of supporting emergency services. For instance, a paper by Koufi et al. [7] describes an application that provides patients with the means to gather critical health information in one place. Specifically, they can print a unique emergency card providing information on how to access these records in case of an emergency. Another example is the Emergency Care Summary (ECS) system implemented by the Scottish Government [4], which provides valuable information about patients to healthcare staff, accessible only under express patient consent. However, in both cases, the summary does not necessarily include detailed medical events. For example, recent diseases, surgeries or disabilities are not part of the summary.

Health records include details about medical events, information that could be used to identify vulnerable people, or people otherwise requiring special assistance, in the context of an emergency. Let's consider, for example, the case of a fire emergency in an academic building. Among the people in the building, there are two who may require special assistance. One person is a wheelchair user, the other suffers from lung disease. The wheelchair user may obviously need support to evacuate, however, such information is not known to emergency services. Crucially, the person's health records include a diagnosis of a vertebral fracture (a permanent condition). An intelligent system with access to this information would understand that such condition does not improve over time. In this way, emergency services may be alerted to prioritize effort in supporting this person. However, the second case is more subtle. The convalescence process for "obstructive bronchitis" means this person has difficulty breathing and walking. Typically, if the issue is recent, symptoms are rather acute. However, the health of the person improves with appropriate treatment, until full recovery. An intelligent system can analyse his/her health records and provide valuable information about their permanent or recent health issues and how severe they are at the time of the emergency.

In [11], we proposed a four steps methodology to support the reuse of health records for identifying vulnerable people in the context of an emergency. We applied a knowledge engineering approach to identify vulnerable people and used a common-sense knowledge base to categorise health records into a type of disability. We implemented a system based on the proposed methodology and made use of synthetic health records to evaluate its efficacy. In order to support the identification of relevant data points, we built a manually curated database of time validity annotations, which

consisted of a sample of SNOMED CT¹ concepts [3]. These health conditions were manually annotated with a *time validity*, which meant that a condition may be valid for a specific amount of time. Our system used these annotated health records to decide whether a person required assistance. However, the approach presented in [11] has some limitations:

- Time validity annotations are expressed in a simplistic way (number of months required to heal). As a result, the system includes many irrelevant health records.
- The current representation of time validity falls short of specificity and, therefore, limits recommendations about the severity of a condition. For instance, the model annotates a health condition as chronic or temporal, thus missing cases where conditions deteriorate over time or improve after a convalescence period.
- Current time validity annotations only cover a small number of SNOMED CT concepts. Health records use a wide variety of SNOMED CT concepts to describe medical events. Collecting time validity information becomes essential for implementing the approach.
- Although in [11] results report satisfactory accuracy and recall, precision (the number of people identified as vulnerable that were correctly classified) leaves room for improvement.

The process of recovering from a health situation is not limited to stating the convalescence time. There are situations such as chronic conditions that could deteriorate over time. Crucially, events such as symptoms and procedures in health records refer to conditions typically described as having a minimum and maximum recovery time.

In this paper, we propose a methodology for representing and reasoning on condition evolution information. We introduce the notion of *Condition Evolution Statement (CES)* to solve the problem of representing and reasoning over the health status of a person at a given point in time. This notion is different from the one of time validity, introduced in [11]. *Time validity* annotations represent a fixed recovery time or a chronic condition. Instead, *Condition Evolution Statement* is an elaborated representation of the recovery process, which defines the duration of the convalescence process and the type of issue. Therefore, it makes it possible to identify if the condition is temporal or chronic and its progression. In other words, how severe it is at a given point in time. To the best of our knowledge, there is no previous work on the representation and reasoning of condition evolution for supporting the detection of vulnerable people in emergency response.

Similarly, there is no existing structured data about condition evolution available for reuse. Therefore, we present an approach to building such a detailed database of condition evolution. We rely on knowledge acquisition techniques such as Machine Learning (ML) to build a semi-automatic supervised classification pipeline. The approach uses unstructured text collected from public health websites, such as NHS² and MAYO³. We train different ML algorithms and classify sentences according to the different components of the

CES representation. Ultimately, we put together the results of the classification and generate a condition evolution statement.

Finally, we reproduced the same evaluation setting of [11]. We replaced the system's element 'time validity' with the two new components: the database and a model based on CES. We ran extensive experiments, which show how the system improved its precision by four points. The results prove that our Intelligent System enhances the detection of vulnerable people by making it more accurate thanks to the representation of a condition's evolution over time. Furthermore, the CES representation now allows us to evaluate the severity of health issues, a new feature providing more detailed information to emergency services.

Our main contributions⁴ are:

- A model to represent and reason on condition evolution, the Condition Evolution Statement (CES).
- A pipeline that allows the *semi-automatic* extraction of information about condition evolution.
- A dataset of 1324 SNOMED CT concepts annotated with a condition evolution statement. To the best of our knowledge, our dataset is the first to annotate SNOMED CT and provide detailed information about condition evolution.
- A training dataset with 1987 sentences that describe the recovery time of different conditions. We use this dataset to train the ML algorithms on the text classification task.
- An extensive set of experiments using a synthetic dataset of health records demonstrates that our new approach enhances the precision of our system to distinguish vulnerable people. Furthermore, it provides means to know the probability that an individual requires assistance with a severity score.

The rest of the paper is organised as follows. In Section 2, we present the proposed methodology. In Section 3, we describe the approach to building the dataset and the CES representation. Section 4 reports the results of the different experiments and the comparison of the results against previous work. In Section 5, we review the literature on the use of health records for emergency. Finally, in Section 6 we summarise the main conclusions and outline future directions of research.

2 METHODOLOGY

Section 1 described the limitations associated with using time validity representation to detect vulnerable people in an emergency. A more sophisticated description of how health conditions develop over time can enhance the system's precision, detect ongoing conditions and provide means to calculate the recovery duration. Here, we propose a methodology for designing a model to represent condition evolution, allowing the inference of ongoing health issues. We complement the proposed methodology by constructing a database of SNOMED CT concepts, augmented with CES annotations. In what follows, we summarise the phases of our methodology.

Data collection. The first phase of the methodology focuses on identifying information sources about condition evolution. Such resources are necessary for two reasons: (a) develop a knowledge model about health conditions' evolution; and (b) populating a

¹Systematized Nomenclature of Medicine, Clinical Terms. Owned by SNOMED International, it is a comprehensive logic-based clinical healthcare terminology used to present clinical content in electronic health records.

²<https://www.nhs.uk/conditions/>

³<https://www.mayoclinic.org/diseases-conditions>

⁴The database and the Python based code is available here: <https://github.com/albamoralest/CES-HR>

database of conditions’ evolution. The sources should comply with requirements such as: being publicly available, extensive, including a description of condition evolution and authoritative (reliable sources). The aim is to collect text descriptions of diseases, procedures, conditions (for instance, asthma, appendicitis, bronchitis). In this paper, we refer to them as health conditions or simply conditions.

Knowledge Representation. Next, our attention is given to representing and reasoning on condition evolution. First, we examine how condition evolution is expressed in natural language from the selected sources. The objective is to identify the different features that define a condition evolution. For instance, the recovery time generally expresses the direction (whether a condition improves or declines), and a time range, i.e how long it takes to recover (for instance, from 1 week to 1 month, in 3 weeks, one year, etc...). With this information, we design a model for representing the evolution of health conditions, which comprises a set of Condition Evolution Statements (CES).

Knowledge Acquisition. In this phase, we focus on the knowledge acquisition process, using the data sources identified and the CES model. The objective is to populate a database of CESs, linked to SNOMED CT. The process is divided into a number of steps.

- (1) We implement a Machine Learning pipeline that automatically classifies sentences according to the CES model. However, the CES has different elements; therefore, we train a set of models [10], each specialising in identifying one component. The output is a collection of health records with sentences annotated with condition evolution statements.
- (2) Because a condition can have one or many sentences annotated with a CES; thus, the next task is to select the best CES for each health condition. We apply an algorithm that uses support and confidence as metrics to identify the most frequent combination of annotations and blends the results to generate a CES for each sentence.
- (3) The last task is to match each health condition with its corresponding SNOMED CT concept. The output of this phase is a dataset of SNOMED CT concepts, each one of them linked to a condition evolution statement.
- (4) However, the information sources identified in the first phase have a limited coverage of SNOMED CT concepts. Therefore, we exploit the semantic structure of SNOMED CT to derive propagation rules and expand the coverage of the CES database to other concepts in SNOMED CT. Applying the propagation rules makes it possible to elicit an extensive and structured dataset of SNOMED CT concepts with CES annotation.

Reasoning on condition evolution. Finally, we study how a system can use CES to predict whether a condition diagnosed in a health record is still ongoing at a certain point in time. Specifically, a health record is relevant if the condition is still ongoing (an individual has not recovered from it) when the emergency occurs (e.g. the fire started in the building). For example, a condition that generally improves in two weeks and occurs three days before the fire event is likely to impact a person’s health, who might require assistance to evacuate. The system will use the CES to calculate if the condition is ongoing and, therefore, valid.

3 APPROACH

This section describes how we applied our methodology to represent and reason with conditions’ evolution and construct a database of SNOMED CT concepts annotated with CESs.

3.1 Data collection

As described in the methodology, the first task is to identify data sources from (a) an authoritative organisation and (b) publicly available. Also, sources should be (c) extensive and (d) contain descriptions of condition evolution. Here we rely on two health organisations: NHS England and MAYO Clinic. NHS England is the largest health website in the UK, and it provides straightforward access content about symptoms, conditions, and treatments. The MAYO Clinic is a non-profit organisation oriented to clinical practice, education, and research, providing comprehensive and easy access to condition descriptions. NHS England website displays information of 972 health conditions and MAYO Clinic, 1170 health conditions. Both include sections that describe the ‘recovery’ and ‘treatment’ where we can find condition evolution information.

From both websites, we collect HTML files that contain health conditions descriptions expressed in natural language. We clean the text by removing HTML tags, line breaks, special characters and empty spaces. After reviewing the text descriptions, we noticed that condition evolution is usually described in one sentence. For instance, “Bronchitis” web page contains sentences such as: *“In most cases, acute bronchitis clears up by itself within a few weeks without the need for treatment.”* and *“If symptoms last for at least 3 months, it’s known as chronic bronchitis.”* We organised the text by sentences and removed sections such as “Preparing for your appointment” and “Sources”. Our final dataset has 214,012 sentences in total, grouped by health conditions.

3.2 Knowledge representation

Here, we focus on building a representation of condition evolution that can be used for a system. First, we use a subset of sentences describing condition evolution from the outcome in the previous step. In order to create this subset, we manually select text snippets indicating condition recovery, for instance: *fully recover, last between, is a progressive condition/disease, lifelong condition, no specific cure*, among others. Next, we use this list and cosine similarity measure to find complete sentences with condition evolution.

We use this subset of sentences to analyse its structure and abstract three dimensions: direction, pace, and time range. For each dimension, we establish annotations to represent all the possible values they could take. We do this by grouping expressions with similar meanings. Table 1 lists the annotations of our model, their interpretation, and the list of *expressions* found in these sentences. For example, the direction feature indicates if a condition gets better or is chronic. Then, certain expressions also give details about the speed of the recovery. And finally, a measurable account of how long it last (days, weeks, months, years). Figure 1 displays the structure of our model.

By adopting the CES Model, a sentence or a piece of text describing a condition evolution can be represented in a machine-readable way as a combination of these features. However, not all combinations of features are meaningful. For example, the features NONE

Table 1: Summary of expressions used to develop the Condition Evolution Statement annotations

Dimension	Annotation	Annotation definition	Expressions found in sentences
DIRECTION	IMPROVE	Indicates recovery of good health condition. Always have a minimum and maximum recovery time	"improve in", "no more than X minutes", "last between/within/around X", "take around", "less than", "fully recover", "temporary", "usually doesn't need treatment", scans and tests with specific start and end date.
	DECLINE	Indicates that gradually becomes worse.	"deteriorates", disorder "develop slowly/rapidly/gradually", "is a progressive condition/disease", "gets gradually worse over time", progression in short/long periods of time, a developing condition that lead to death, that cause disability.
	PERMANENT	Long lasting and never goes away.	"lifelong condition", "no specific cure", "cannot be cured", "it is a long-term condition/complication", not regaining full pre-injury status, lifelong conditions that affect everyday life, when treatment is needed for the rest of life.
	NONE	Describe administrative procedures, not affecting health.	Scans and tests, do not specify start and end date because are very quick and do not cause problems.
PACE	FAST	That happens very quickly.	"rapidly", "in less than a few/X days", "is a straightforward process", "very quick".
	MODERATELY	A health condition that progress as expected.	"develop gradually", "several weeks", "within a few months", recovery time falls between fast and slowly.
	SLOWLY	A health condition happens in a long period.	"develop slowly", "several months to years", "over years/many years/several years", "often over several months", "long time to recover (from few months to a year)", "progresses slowly", long time to recover, recover slowly.
TIME RANGE	FROM	Minimum period of convalescence or deterioration starts, also called lower bound (LB).	Expressed in hours, days, months or years, several/few years/months/weeks.
	TO	Maximum period of convalescence or deterioration becomes a chronic/permanent condition, also called upper bound (UB).	Expressed in hours, days, months or years, several/few years/months/weeks.

Figure 1: Condition Evolution Statement representation

Table 2: Examples of sentences representation using CES

Sentence	Condition Evolution Statement (CES) Direction + Pace + Time
Addison's disease symptoms usually develop slowly, often over several months.	DECLINE SLOWLY FROM 6 MONTHS TO 1 YEAR
A broken ankle usually takes 6 to 8 weeks to heal, but it can take longer.	IMPROVE MODERATELY FROM 8 DAYS TO 2 MONTHS
There's currently no cure for CJD, so treatment aims to relieve symptoms and make the affected person feel as comfortable as possible.	PERMANENT
Fast treatment is essential	NONE

and PERMANENT express either the absence of a condition evolution expression or a condition evolution expression that, although relevant, does not change over time. Therefore, NONE and PERMANENT do not combine with "PACE" and "TIME RANGE". Besides, the combination of pace and time range features should be coherent. For example, if the time range expression is "FROM 2 MONTHS TO 6 MONTHS", it cannot be "FAST". In summary, the sentence "Most people will fully heal within 2 to 8 weeks" becomes: "IMPROVE MODERATELY FROM 8 DAYS TO 2 MONTHS". Table 2 displays few examples of sentences and their reformulation as CES.

3.3 Knowledge Acquisition

In this section, we describe how we populated a database of CESs. We develop a knowledge engineering approach and build a semi-automatic supervised classification pipeline to extract condition evolution information from natural language sources (NHS and MAYO conditions descriptions).

3.3.1 Sentences classification. Here the aim is to identify the sentences that indicate health condition evolution from health condition web pages extracted in the data collection phase. We apply machine learning to automatically classify them according to the dimensions used in CESs: direction, pace, and time range. The output consists of a set of sentences for each condition with the corresponding CES. In what follows, we describe the process to build a database of health conditions and CES.

Building the training dataset. In Section 3.2 we built an initial sample of sentences describing condition evolution. Using the CES representation, we assign the corresponding annotation manually to each sentence. In anticipation of the ML classification task, we complete the training sample by adding negative annotations, this means, sentences without a description of condition evolution. The output is a manually curated training set of 1987 sentences and their CES. Table 3 summarises the total number of sentences grouped by CES.

Experiments with ML algorithms. The first objective is to discriminate sentences that describe condition evolution from those that do not contain such description. We add a step (C0), applying a boolean classifier aimed at distinguishing sentences that include a condition evolution expression, from others that do not. As a result, we obtain a reduced dataset of sentences which have a good chance of containing condition evolution information. As described in Subsection 3.2 the CES is represented using three different features. Therefore, we classify a sentence across the three dimensions of the CES. We train different Machine Learning (ML) models for

Table 3: # of sentences per CES in the training dataset

Condition Evolution Statement (CES)	Total
NONE	1437
PERMANENT	141
IMPROVE MODERATELY FROM 8 DAYS TO 2 MONTHS	106
IMPROVE FAST FROM 5 MINUTES TO 1 DAY	74
DECLINE SLOWLY FROM 1 YEAR TO MORE YEARS	56
IMPROVE MODERATELY FROM 2 MONTHS TO 6 MONTHS	53
IMPROVE FAST FROM 1 DAYS TO 1 WEEKS	37
IMPROVE SLOWLY FROM 1 YEAR TO MORE YEARS	37
IMPROVE SLOWLY FROM 6 MONTHS TO 1 YEAR	30
DECLINE FAST FROM 1 DAY TO 1 WEEK	6
DECLINE SLOWLY FROM 6 MONTHS TO 1 YEAR	4
DECLINE MODERATELY FROM 8 DAYS TO 2 MONTHS	4
DECLINE MODERATELY FROM 2 MONTHS TO 6 MONTHS	2
TOTAL sentences	1987

Table 4: ML training results: Accuracy for each feature and different models

ML Algorithms	C0	Direction	Pace	Time range
<i>Logistic Regression</i>	0.8907	0.9727	0.8148	0.8114
<i>Decision Tree</i>	0.8337	0.9272	0.8934	0.8606
Linear SVC	0.8789	0.9545	0.8271	0.8360
MLP Classifier	0.8337	0.9545	0.7530	0.6803
Naïve Bayes	0.4181	-	-	-
Multinomial NB	0.8136	0.8636	0.6790	0.5737
Random Forest Classifier	-	0.8818	0.7901	0.8442

each CES feature and obtain a piece of what could be the CES. We randomly divide the dataset into a separate training (70%) and test (30%) dataset. Table 4 shows the list of the ML algorithms we trained and the accuracy of each model.

Application of the machine learning approach. From the previous step, we use the ML models with the best performance for each feature of the CES representation. A total of 5174 out of 214,012 sentences were classified as providing information about condition evolution and assigned a CES. A total of 3635 different sentences were selected, once we deleted the repeated combinations of "sentence+CES" within a given condition. A total of 1439 health conditions have one or many sentences with a CES annotation.

3.3.2 Consistency check and blending. The result from the previous step is a dataset of sentences with a CES annotation. However, a health condition can have one or more sentences with a CES. Thus, we apply an algorithm to verify which combinations of feature annotations are coherent, for this we use the PACE feature as described in Subsection 3.2 to review consistency. For instance, we discard combinations such as 'FAST' and 'FROM 6 MONTHS TO 1 YEAR', or the contrary 'SLOWLY' and 'FROM 5 MINUTES TO 1 DAY'.

Next, we use an association rule learning method to identify how likely it is for a combination of CES features to represent a health condition. Firstly, we calculate how frequently the combination of *health condition and CES (direction+pace+time range)* appears in the dataset (support). Then we calculate how often the combination of *health condition and CES is valid* (confidence). Finally, we select the CES with the best confidence, which represents a health condition. In Table 5 we exemplify the use of these metrics to select the CES.

Finally, we link the health conditions with its equivalent SNOMED CT concept. As a result, a total of 1,324 SNOMED CT concepts now have a CES annotation.

3.3.3 Knowledge propagation. Using an ML pipeline, we managed to annotate a total of 1,324 SNOMED CT concepts. However, the number of annotations corresponding to the concepts in the initial sample, the HR sample (used for experiments of our previous work) is 16%. In order to cover the whole sample, we analysed the SNOMED CT taxonomy, looking for shared attributes that could indicate that two or more concepts share a CES. Specifically, we looked for shared features that indicate similarity or other relationship between concepts in SNOMED CT. The objective is to discover patterns that could guide CES's propagation, from a concept with CES to another without CES.

In total, six rules were created and applied to the SNOMED CT concepts. They covered 5000 concepts in total. Using the propagation rules, we managed to annotate 91% of the original HR sample. Our objective is to reproduce the experiments from our previous work, and therefore for the remaining concepts, we chose to annotate them manually.

3.3.4 Manual supervision and validation. Finally, we validated the output of the process. We compared the output against the database used in the previous version of our system [11]. The results show that 54% of concepts are consistent with the time validity database (manually annotated). We reviewed the remaining annotations and realised that 60% of these annotations are better expressed with the new CES. The knowledge acquisition process proved helpful to build the new database and running the experiments with the CES representation. The evaluation of the knowledge acquisition process is beyond the scope of this paper, however, is considered for future work.

3.4 Reasoning with Condition Evolution Statements

In this section, we explain how a system can use CESs to predict whether a certain condition holds at the time of an emergency.

Typically, a condition recorded in a health record is relevant if it is still ongoing. This happens for three reasons: recovery time has not passed yet, the condition is chronic, or the condition deteriorates in time. We use the direction feature with two purposes: a) to identify the *relevant* conditions and b) to guide the reasoning on time range. In what follows, a description of the logic involved for each annotation is given:

- (1) NONE: If the direction's annotation is "NONE" then the system marked this health record as NOT RELEVANT.
- (2) PERMANENT: If a condition has as annotation PERMANENT, then the condition is always valid.
- (3) IMPROVE: Describes a health condition that lasts for a certain amount of time. The CES represents the convalescence process using the annotations FROM and TO. Typically, the convalescence period could last a minimum time (FROM) or lower bound (LB) in the best-case scenario. In the worst-case scenario, a maximum time (TO) or upper bound (UB). (see Figure 2a)

Table 5: Dataset results: Best confidence

SNOMED Concept	SNOMED Identifier	Condition Name	CES	Confidence	Sentence	Source
Abdominal aortic aneurysm (disorder)	233985008	Abdominal aneurysm	IMPROVE MODERATELY FROM 8 DAYS TO 2 MONTHS	0.0032	Full recovery is likely to take a month or more.	MAYO
Abdominal aortic aneurysm screening (procedure)	698356002	Abdominal aneurysm screening	IMPROVE FAST FROM 5 MINUTES TO 1 DAY	0.0020	Screening for AAA involves a quick and painless ultrasound scan of your tummy.	NHS
Chronic obstructive lung disease (disorder)	13645005	COPD	DECLINE SLOWLY FROM 1 YEAR TO MORE YEARS	0.0036	Although COPD is a progressive disease that gets worse over time, COPD is treatable.	MAYO
Chronic obstructive lung disease (disorder)	13645005	Chronic obstructive pulmonary disease (COPD)	PERMANENT	0.0034	There's currently no cure for chronic obstructive pulmonary disease (COPD), but treatment can help slow the progression of the condition and control the symptoms.	NHS
Cyst of ovary (disorder)	79883001	Ovarian cyst	IMPROVE SLOWLY FROM 2 MONTHS TO 6 MONTHS	0.0082	In most cases, ovarian cysts disappear in a few months without the need for treatment.	NHS
Cyst of ovary (disorder)	79883001	Ovarian cysts	IMPROVE SLOWLY FROM 2 MONTHS TO 6 MONTHS	0.0082	The majority disappears without treatment within a few months.	MAYO

- (4) **DECLINE**: This annotation describes a condition that deteriorates over time. For instance, a person that has good health but showing certain symptoms. At some point, the person is diagnosed (a health record is created: RD). He/she receives care, but due to the nature of the condition, the person's health continues getting worse (FROM/LB). After some time it becomes chronic or permanent (TO/UB) (see Figure 2b).

Figure 2: Reasoning on CES: Direction interpretation.

After evaluating the direction feature in CES, we move to reason on TIME RANGE and how a system should interpret it in combination with CES feature direction. As stated in Table 1, the TIME RANGE annotation has two elements: a lower bound (LB) value and an upper bound (UB) value. Whether a health record is valid or not depends on whether the emergency is happening before, in between or after these two boundaries.

- (1) The emergency happens after the UB. On the one hand, if direction annotation is IMPROVE, then the convalescence period has ended (see Figure 3a). Thus the health record is not relevant. On the other hand, if direction is DECLINE, then the condition has become permanent, therefore the health record indicates a current health issue (see Figure 3b). Hence, it is relevant.
- (2) The emergency happens in between LB and UB: the health record is always relevant.
- (3) The emergency happens before the LB. In the case of IMPROVE annotation, the condition is valid as the person is still suffering from the given condition (see Figure 3c). In the

case of DECLINE, it is possible that the person is showing symptoms related with a condition, but the deterioration process starts after the LB (see Figure 3d).

Figure 3: Reasoning on CES: Time range interpretation.

3.4.1 Severity score. The CES representation provides the opportunity to assess if a person is recently recovering from a health condition and, therefore, likely to require assistance. For instance, if a condition improves with time, the closer the emergency date is to the minimum recovery time, the higher the probability of requiring assistance is (see Figure 4a). On the contrary, if the condition declines, the closer the emergency date is to the minimum time for the condition to start deteriorating, the lower the probability of requiring assistance (see Figure 4b). If a condition is PERMANENT, the condition is ongoing, therefore, severity is always the highest.

In order to calculate the *severity score*, we establish a six-level system. Level one represents the period before LB, levels two to five the quartiles between LB and UB, and finally level 6 the period after UB. The system calculates the severity score according to the proximity to the "Poor health" zone and assigns the severity score accordingly. Figure 4 summarises the use of quartiles for severity score calculation.

Figure 4: Reasoning on CES: Severity score.

4 EVALUATION

Our objective is to enhance the precision of a system that uses health records to identify vulnerable people during a fire emergency. One way of achieving this objective is to improve the classification of timely valid health records that could reveal information about health issues affecting people's ability to evacuate. In this paper, we introduced a methodology that allows us precisely to represent and reason on condition evolution. To complement this work, we also built a database with information about condition evolution.

We replicate the experiment setting of our previous work [11]. We use the same 1012 patients' health records sample. Similarly, to evaluate the results of our experiments, we use the gold standard developed in our previous work. In these experiments, we replace our system's previous time validity component with the outcome from implementing our approach—the two new components: the CES reasoning engine and the database of SNOMED concepts with condition evolution statements.

4.1 Experiments

While in our previous work, the time validity was a simplistic representation in months. The current CES model establishes a detailed and precise description of condition evolution. In order to make the best use of the CES representation, we designed three experiments, each one using the time range feature in different ways. In what follows, we describe each experiment considering various hypotheses.

4.1.1 Experiment one - Pessimistic approach or upper bound (UB). For this experiment, the assumption is that a condition develops in the most prolonged period. For example, suppose a condition improves around two days (LB) to 8 days (UB). In that case, the system assumes the disease takes the longest time to recover (8 days). Therefore, if the emergency happens in this period (including the 8th day), the condition is valid, and the person might need assistance.

4.1.2 Experiment two - Optimistic approach or lower bound (LB). For this experiment, the hypothesis is that a condition develops in the shortest period. Following the previous example, if a condition is meant to improve in 2 (LB) to 8 days (UB), the system assumes that the condition takes the shortest time to recover (2 days). Therefore, if an emergency happens before the LB (the period it takes

Table 6: Experiments results

Experiments	Accuracy	Precision	Recall	F-score
Upper bound (UB)	0.92	0.73	0.81	0.77
Lower bound (LB)	0.90	0.79	0.51	0.62
Average (Av)	0.92	0.80	0.72	0.76
Simplistic approach - [11]	0.91	0.69	0.82	0.75

to recover: two days), the condition is valid. If the emergency happens between days 3 and 8, the person is assumed to have already recovered his/her good health, thus not requiring help.

4.1.3 Experiment three - Average or median approach (Av). This experiment assumes that a health condition develops in a period considered the median between the LB and UB. Continuing with the previous example, the median number of days is 3 days. In that case, the system assumes the person's convalescence period finishes on 5 days. If the emergency happens before that day, then he/she might still require help.

4.2 Results

We incorporated the outcome of the knowledge acquisition and representation approaches and compared the results of our system against the gold standard generated in our previous work [11]. The results obtained from this analysis are summarised in Table 6, which also includes a row with the results of the previous work "Simplistic approach - [11]".

As seen in Table 6, in the three experiments, the overall accuracy of the system remains equivalent. Crucially, our approach reported clear evidence of an increase in Precision in all the experiments, particularly in experiment three (Av), the average approach.

We also observe that improving Precision can affect Recall, as shown in experiment LB - lower bound. Experiment one, *Pessimistic approach (UB)*, provides the best result for Recall. Results report a 4 point increase in Precision, while maintaining Recall's performance. Since the general objective of the system is to report on people who need assistance, Experiment one provides the best balance between precision and recall. However, to improve the quality of the recommendations, our approach enables the system to additionally provide a *severity score*, to leverage the predictions of LB and AV, having higher precision.

Following examples from Section 1, Table 7 shows the severity score for the wheelchair user and the person with lung disease. In both cases, the severity score is one, the highest score. The "vertebral fracture" indicates a PERMANENT condition. At the same time, "chronic bronchitis" improves in a certain period. However, as it started recently (considering a fire emergency 25 days later), the system assumes the person has some difficulties as the recovery period has not finished.

5 RELATED WORK

This section reviews related work about the representation of health records and medical terms, use of health records to support emergency services and extraction of medical information from natural language.

Table 7: Examples severity score

Condition	Date	Quartile	Severity score	CES annotation
Chronic obstructive bronchitis (disorder)	05/10/2019	LB	1	IMPROVE MODERATELY FROM 2 MONTHS TO 6 MONTHS
Fracture of the vertebral column with spinal cord injury	08/04/2008	LB	1	PERMANENT

The main benefit of having patients’ health records in digital format is that they can be easily collected, processed, and accessed. Standardised terminology for recording clinical content such as SNOMED CT facilitates these tasks [1, 8]. SNOMED CT, which has been widely adopted (currently 41 member countries, more than 5000 licences issued), includes 352,567 clinical terms, grouped in 19 domains in a polyhierarchical structure. The core components of SNOMED CT are concepts, descriptions and relationships that enable accurate capture and representation of patient’s medical information [3]. As a testbed for evaluating the proposed approach, our system uses a synthetic dataset of health records. Synthea [12] is an open-source software that models a deep and extensive medical history of patients. Furthermore, it employs SNOMED CT terminology to generate health records. As healthcare records are increasingly digitised, FHIR (Fast Healthcare Interoperability Resources) [5] is the standard specification adopted to represent the health records dataset.

Temporal information of health records could reveal medical issues that serve the purpose, for example, to identify disabilities. As reported by [2, 9], significant research has been conducted to extract medical information from natural language sources. To the best of our knowledge, none of these approaches studies the representation of the evolution of health conditions, in the context of supporting emergency services. Furthermore, there is no structured information about health conditions recovery.

The increased availability of digitally collected health records has facilitated its use for emergency services. One example is monitoring healthcare data changes in lone older people, notifying first responders when an emergency happens [6]. Similarly, health records have been used to construct a tailored emergency card [4, 7] that contains patient’s health information (blood type, allergies). While it may contain important information, the card does not include medical events (recent surgeries, diseases) to identify vulnerable people. In our previous work [11], we approached the problem of identifying vulnerable people in the context of a fire evacuation by implementing a pipeline that makes use of healthcare data. The system classifies people that require assistance by identifying time relevant health records. Although the results of our previous work proved a reasonable good accuracy to detect people in need, we identified important limitations of the approach. Specifically, our previous system precision could improve by utilising more accurate values to express recovery time. We identified that the time validity representation used by the previous system falls short of specificity and does not provide the means to calculate the severity of the health issue.

6 CONCLUSIONS

In this paper, we proposed a methodology for representing and reasoning about the evolution of health conditions in the context of supporting emergency services. We introduced the concept of Condition Evolution Statement (CES) and constructed the reasoning behind this representation. Furthermore, we built a database of condition evolution annotations to support our model and demonstrated that it significantly improves the quality and precision of the system.

Future work includes evaluating the condition evolution annotations produced as a result of the knowledge acquisition phase of the methodology. As described in this paper, we used authoritative, extensive and open sources that describe health conditions evolution. However, providing means to review and validate the methodology and data by domain experts is essential and would also contribute to a better coverage of SNOMED CT. Additionally, we consider that our proposed model for representing and reasoning on health condition evolution can be formally encoded using Web Ontology Language (OWL). Following the Linked Data approach, publish the dataset of condition evolution as a structured information resource.

REFERENCES

- [1] Ankur Agrawal and Paul Revelo. 2017. Analysis of the consistency in the structural modeling of SNOMED CT and CORE problem list concepts. In *2017 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*. IEEE, Kansas City, MO, 292–296. <https://doi.org/10.1109/BIBM.2017.8217666>
- [2] Ghada Alfattni, Niels Peek, , and Goran Nenadic. 2020. Extraction of temporal relations from clinical free text: A systematic review of current approaches. *Journal of Biomedical Informatics* 108 (aug 2020), 12 pages. <https://doi.org/10.1016/j.jbi.2020.103488>
- [3] SNOMED CT. 2017. *SNOMED CT Starter Guide*. <http://snomed.org/srg>
- [4] Scottish Government. 2020. *Your Emergency Care Summary: What does it mean for you?* <https://www.gov.scot/publications/emergency-care-summary-mean/>
- [5] HL7.org. 2019. *HL7FHIR*. <https://www.hl7.org/fhir/overview.html>
- [6] Aamir Hussain, Rao Wenbi, Aristides Lopes da Silva, Muhammad Nadher, and Muhammad Mudhish. 2015. Health and emergency-care platform for the elderly and disabled people in the Smart City. *Journal of Systems and Software* 110, 7 (dec 2015), 11 pages. <https://doi.org/10.1016/j.jss.2015.08.041>
- [7] Vassiliki Koufi, Flora Malamateniou, and George Vassiliopoulos. 2010. Automating Authorization Propagation Process In Personal Health Records. In *Proceedings of the Third International Conference on Health Informatics*. SciTePress, Valencia, Spain, 82–89. <https://doi.org/10.5220/0002747200820089>
- [8] Dennis Lee, Nicolette de Keizer, Francis Lau, and Ronald Cornet. 2014. Literature review of SNOMED CT use. *J Am Med Inform Assoc* 21 21, 1 (feb 2014), 8 pages. <https://doi.org/10.1136/amiajnl-2013-001636>
- [9] Fang Li, Jingcheng Du, Yongqun He, Hsing-Yi Song, Mohcine Madkour, Guozheng Rao, Yang Xiang, Yi Luo, Henry W Chen, Sijia Liu, Liwei Wang, Hongfang Liu, Hua Xu, and Cui Tao. 2020. Time event ontology (TEO): to support semantic representation and reasoning of complex temporal relations of clinical events. *Journal of the American Medical Informatics Association* 27, 7 (jul 2020), 11 pages. <https://doi.org/10.1093/jamia/ocaa058>
- [10] Kotsiantis SB, Zaharakis I, and Pintelas P. 2007. Supervised machine learning: A review of classification techniques.. In *Emerging artificial intelligence applications in computer engineering*. IOS Press,US (1 Oct. 2007), 3–24. https://www.enib.fr/~chevaill/documents/master/siia_bibl/Kotsiantis_i07.pdf
- [11] Alba Catalina Morales Tirado, Enrico Daga, and Enrico Motta. 2020. Effective Use of Personal Health Records to Support Emergency Services. In *Knowledge Engineering and Knowledge Management*, C. Maria Keet and Michel Dumontier (Eds.). Springer International Publishing, 54–70. https://doi.org/10.1007/978-3-030-61244-3_4
- [12] Jason Walonoski, Mark Kramer, Joseph Nichols, Andre Quina, Chris Moesel, Dylan Hall, Carlton Duffett, Kudakwashe Dube, Thomas Gallagher, and Scott McLachlan. 2018. Synthea: An approach, method, and software mechanism for generating synthetic patients and the synthetic electronic health care record. *Journal of the American Medical Informatics Association* 25, 3 (mar 2018), 8 pages. <https://doi.org/10.1093/jamia/ocx079>