ARTICLE IN PRESS

Artificial Intelligence In Medicine xxx (xxxx) xxx–xxx

Contents lists available at ScienceDirect

Artificial Intelligence In Medicine

journal homepage: www.elsevier.com/locate/artmed

Using preference learning for detecting inconsistencies in clinical practice guidelines: Methods and application to antibiotherapy

Rosy Tsopra[a,b,1,*], Jean-Baptiste Lamy[a,1], Karima Sedki[a]

[a] LIMICS, INSERM UMRS 1142, Université Paris 13, Sorbonne Université, F-75006 Paris, France
[b] AP-HP, Paris, France

ARTICLE INFO

Keywords:
Preference learning
Antibiotherapy
Clinical practice guidelines
Inconsistencies in guidelines

ABSTRACT

Clinical practice guidelines provide evidence-based recommendations. However, many problems are reported, such as contradictions and inconsistencies. For example, guidelines recommend sulfamethoxazole/trimethoprim in child sinusitis, but they also state that there is a high bacteria resistance in this context. In this paper, we propose a method for the semi-automatic detection of inconsistencies in guidelines using preference learning, and we apply this method to antibiotherapy in primary care. The preference model was learned from the recommendations and from a knowledge base describing the domain.

We successfully built a generic model suitable for all infectious diseases and patient profiles. This model includes both preferences and necessary features. It allowed the detection of 106 candidate inconsistencies which were analyzed by a medical expert. 55 inconsistencies were validated. We showed that therapeutic strategies of guidelines in antibiotherapy can be formalized by a preference model. In conclusion, we proposed an original approach, based on preferences, for modeling clinical guidelines. This model could be used in future clinical decision support systems for helping physicians to prescribe antibiotics.

1. Introduction

In the 1990s, the concept of *Evidence-Based Medicine* was introduced and defined as “the integration of best research evidence with clinical expertise and patient values” [1]. This new paradigm led to the redaction and diffusion of Clinical Practice Guidelines (CPGs) by national health authorities [2]. CPGs are narrative documents providing recommendations stated by a group of experts according to a systematic review of the available clinical evidence. They aim at improving the quality of health care by providing standardized best practices for diagnosis and treatment. Their development is complex and requires time, rigor and multiple verification and validation steps to guarantee their quality [3–6]. However, many problems are reported like incompleteness, contradiction, inconsistency, redundancy or ambiguity within CPGs [4]. This leads to a lack of confidence of physicians in CPGs [7], and thus a poor consideration of CPG recommendations in their daily routine clinical practice [8].

For verifying the quality of recommendations within CPGs, various methods were developed. The structure of CPGs can be verified by tools [9,10] such as AGREE instrument [11]. These tools focus on quality criteria, e.g. presentation of guidelines, or independence of experts [12]. However, these methods are limited to the verification of the structure of CPGs, and do not consider the consistency and medical pertinence of recommendations.

The consistency of recommendations can be verified by formal methods [13]. The recommendations are first represented using an explicit and non-ambiguous model in a formal language. Several Computed Interpretable Guidelines (CIG) were developed [14]. They allow detecting ambiguity, incompleteness, inconsistency or redundancy within CPGs [6,13,3,15–17]. For example, some authors [18,19] state that, if narrative guidelines are encoded into logical language (“if… then…” rules), then the generation of all possible variable combinations allows the detection of incompleteness (i.e. variable combinations not covered by CPGs) and inconsistencies (i.e. similar variable combinations leading to different conclusions). But these methods are time-consuming and dependent on the formal language. Moreover, these formal approaches don’t verify the medical pertinence (e.g. they do not verify that the recommended drug treatments are not contraindicated for the patient).

Few approaches have been proposed for verifying the medical pertinence of recommendations. These approaches require the formalization of the medical knowledge involved (e.g. drug properties such as

* Corresponding author at: AP-HP, LIMICS, INSERM UMRS 1142, Université Paris 13, Sorbonne Université, Paris, France.
E-mail addresses: rosy.tsopra@aphp.fr (R. Tsopra), jibalamy@free.fr (J.-B. Lamy), sedkikarima@yahoo.fr (K. Sedki).
[1] These two authors contributed equally to this work.

https://doi.org/10.1016/j.artmed.2018.04.013
Received 12 February 2018; Received in revised form 28 March 2018; Accepted 30 April 2018

contraindications) and the identification of the medical principles underlying the recommendations of CPGs. However, formalizing the knowledge and the reasoning principles is a complex task [13]. For example, in oncology, a medical domain where multiple drugs are often prescribed, the adverse events can be limited by checking the known adverse effects [20].

Recently, many approaches have been proposed for enriching recommendations by integrating additional information. These pieces of information concern particularly patient context (psycho-social, multi morbidity, etc.) and patient preferences [21–26]. For example, in multi-criteria decision making, to recommend an appropriate manual wheelchair, user preferences that are often conflicting must be taken into account [27]. Nevertheless, the manual construction of preferences remains complex and time-consuming. Thus, it is more appealing to learn preferences from data, because in general, data are easily collected or observed.

In this article, we propose a method for the semi-automatic detection of inconsistencies in guidelines using preference learning, and we apply this method to antibiotherapy in primary care. In primary care, CPGs recommend prescribing antibiotics empirically, i.e. without knowing the causative bacterium and its susceptibility to the various antibiotics. The most likely bacteria are guessed from the infectious disease (e.g. cystitis is usually caused by *Escherichia coli*). Then, CPGs recommend an antibiotic according to the various antibiotics features (e.g. susceptibility of the likely causative bacteria, side effects) and the patient profile (e.g. child or adult) [28,29].

In order to detect inconsistencies in these CPGs, we made the following hypotheses: (1) it is possible to learn a preference model from the recommendations and a knowledge base describing the domain; (2) a generic model can be defined for all infectious diseases and all patient profiles encountered in CPGs; and (3) this preference model can be used to detect inconsistencies in CPGs.

The rest of the paper is organized as follows. Section 2 gives background about preference learning, and describes the optimization algorithm we used and the antibiotherapy knowledge base we previously designed. Section 3 describes the preference learning. Section 4 describes the detection of inconsistencies and their validation by a medical expert. Section 5 discusses the methods and the results obtained, and finally, concludes.

2. Background

2.1. State of the art in preference learning

Preferences are basically acquired in two ways: (i) by elicitation from the user (for instance through a sequence of queries/answers) or (ii) by learning them directly from data. Preference elicitation is often time-consuming, especially if the number of alternatives/outcomes is large. Moreover, different elicitation techniques are likely to provide different results. It is then more appealing to learn preference from data which is easy to observe and collect. Preference learning is one of the research problems that have recently received considerable attention in disciplines such as artificial intelligence, machine learning, data mining, decision making and others. It aims to learn and construct a preference model from observed preference information. Once the preference model learned, it can be used for decision making for instance. Preference learning can be formalized within various settings, depending for example on the underlying preference model and the type of input provided to the learning system. We can distinguish three common problems in preference learning [30]: (i) learning from label preferences (also designed as label ranking in the literature because frequently, the predicted preference relation is required to form a total order), (ii) learning from instance preferences (instance ranking) and (iii) learning from object preferences (object ranking). Table 1 summarizes the different ranking problems.

In label preferences problem [31,32], the training data contains a set of instances. A set of pairwise comparisons between labels is associated with each instance, expressing that one label is preferred over another for that instance. The objective is to use these pairwise preferences for predicting a ranking function that attributes for any instance a ranking (a total order in general) of all possible labels. Namely, the task is to rank the set of labels for a new instance (label ranking). Label ranking can be considered as a generalization of the supervised classification problem where an order over class labels is associated with an instance instead of only one class label. As an example of a label ranking problem, consider a set of labels $\mathscr{L}$ representing three types of activities: football, tennis and basket. The training data contains a set of students who have to give a list of pairwise preferences between activities (e.g. $\{(Adrien,\ [football \succ tennis]),\ (Marie,\ [tennis \succ football]\}$. Thus, the aim is to compute a ranking over the labels for each instance. For example, the possible prediction of the learnt function for a student x is football ≻ basket ≻ tennis.

In the setting of learning from instance preferences problem [33], the input contains a set of ordered labels and a set of instances, each one associated with a label. The objective is to find a ranking function that allows ranking a given new set of instances. In case where there are two ordered labels, the problem of learning is often called bipartite ranking problem [34]. In case where there are more than two ordered labels, the problem of learning is often called multipartite-ranking problem [35].

Concerning learning from objects [36,37], the objective is to learn a model that allows determining which object is preferred to another. The training data is given in the form of pairwise comparisons between objects. For this type of learning problems, there is no supervision since no class label is associated with an object and each object is not necessarily represented by a set of features or attributes. As an example, to rank query results of a search engine, user clicks on some of the links in the query result and not on others can be exploited to provide training information. Thus, selected pages are preferred over pages that are not clicked.

Two approaches can be distinguished for preference modeling and learning: quantitative and qualitative approaches. Quantitative preferences learning [38,39] consist mainly to learn a utility function on training data. This function assigns a utility degree (or a score) to each alternative (instance, object or label) following the learning problem. For learning problems that are based on qualitative approach [40–44], the objective is to learn a binary preference relation that compares each pairs of alternatives.

When it comes to modeling utility functions, the task is rather more complicated since users may not be used with this formalism and the problem size could be very large. Utility functions (for example, the one a user is supposed to use while making decisions) can be inferred or estimated from past decisions. In [45], this problem is solved by imposing constraints derived from the data over the set of all utility functions. One could go one step further by searching for the optimal utility function given the available constraints. Among first works dealing with deriving utility functions from data, one can mention [46] where the authors aim at extracting reward functions given optimal behaviors in the context of Markov Decision Processes. The main issues dealt with the literature last years concern noise and data inconsistencies and uncertainty, large search spaces and taking into account data sequence, etc. In [47], the authors proposed an approach to learn utility functions allowing to monitor requirements of a dynamically adaptive system. The learned utility functions map at run time monitoring information to a value assessing how well a requirement is satisfied.

Once the preferred model is learned, there is need to measure its quality of prediction. For this, different performance measures can be used such as precision, recall, NDCG (Normalized Discounted Cumulative Gain), etc. In addition, preference learning methods require optimization algorithms. In this study, we will use the Artificial Feeding Birds (AFB) metaheuristics [48,49]. We developed this metaheuristics in previous works, and we describe it in the following section.