Artificial Intelligence in Medicine 55 (2012) 1–11

Contents lists available at SciVerse ScienceDirect

Artificial Intelligence in Medicine

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

Automated interviews on clinical case reports to elicit directed acyclic graphs

Davide Luciani [a,*], Federico M. Stefanini [b,1]

[a] Unità di Ingegneria della Conoscenza Clinica, Laboratorio di Epidemiologia Clinica, Istituto di Ricerche Farmacologiche 'Mario Negri', Via Mario Negri, 1, 20156 Milano, Italy
[b] Dipartimento di Statistica 'G. Parenti', Università degli Studi di Firenze, Viale Morgagni 59, 50134 Firenze, Italy

ARTICLE INFO

Article history:
Received 10 January 2011
Received in revised form 31 October 2011
Accepted 28 November 2011

Keywords:
Knowledge acquisition
Bayesian networks
Directed acyclic graph
Diagnostic reasoning
Problem based learning

ABSTRACT

Objective: Setting up clinical reports within hospital information systems makes it possible to record a variety of clinical presentations. Directed acyclic graphs (Dags) offer a useful way of representing causal relations in clinical problem domains and are at the core of many probabilistic models described in the medical literature, like Bayesian networks. However, medical practitioners are not usually trained to elicit Dag features. Part of the difficulty lies in the application of the concept of direct causality before selecting all the causal variables of interest for a specific patient. We designed an automated interview to tutor medical doctors in the development of Dags to represent their understanding of clinical reports.
Methods: Medical notions were analyzed to find patterns in medical reasoning that can be followed by algorithms supporting the elicitation of causal Dags. Clinical relevance was defined to help formulate only relevant questions by driving an expert's attention towards variables causally related to nodes already inserted in the graph. Key procedural features of the proposed interview are described by four algorithms.
Results: The automated interview comprises questions on medical notions, phrased in medical terms. The first elicitation session produces questions concerning the patient's chief complaints and the outcomes related to diseases serving as diagnostic hypotheses, their observable manifestations and risk factors. The second session focuses on questions that refine the initial causal paths by considering syndromes, dysfunctions, pathogenic anomalies, biases and effect modifiers. A case study concerning a gastro-enterological problem and one dealing with an infected patient illustrate the output produced by the algorithms, depending on the answers provided by the doctor.
Conclusions: The proposed elicitation framework is characterized by strong consistency with medical background and by a progressive introduction of relevant medical topics. Revision and testing of the subjectively elicited Dag is performed by matching the collected answers with the evidence included in accepted sources of biomedical knowledge.

1. Introduction

Bayesian networks (Bns) are possibly the best tools for reproducing clinical reasoning artificially [1–3]. Often, however, the solution of a real medical problem calls for assessment of a large set of variables and relationships, and in fact the increasing amount of detail required to select a treatment is confirmed by the growing trend toward specialization in the medical profession.

In real medical investigations, two main inferential strategies work with individual clinical problems. One concerns patients resembling previously examined ones who were successfully treated. The doctor may apply procedures known to have been at least partially successful in similar cases [4]. However, when the similarities are too weak the doctor cannot follow a case-based reasoning approach. He may then follow causal reasoning, by attempting a pathophysiological interpretation of the patient's internal body state [5]. While clinical knowledge addresses directly the relation between a disease and its clinical manifestations, basic sciences support causal reasoning with the current scientific understanding of biological phenomena [6]. In this way, the doctor tries to reduce the uncertainty about both the patient's health and the impact of potential treatments. This is conditioned in every case by the extent to which medical knowledge fits the case on hand and, if it does, by the doctor's skills in applying it [7].

Here, we consider the causal reasoning of a doctor interviewed about a clinical case report. Case reports contain the description of a patient's clinical presentation and they are widespread in clinical settings, because they are a part of the clinical charts prepared and archived for every hospital patient. They can also be found in

* Corresponding author. Tel.: +39 02 3901 4271; fax: +39 02 39014267.
E-mail addresses: dluciani@marionegri.it (D. Luciani), stefanini@ds.unifi.it (F.M. Stefanini).
[1] Tel.: +39 055 4237266; fax: +39 055 4223560.

0933-3657/$ – see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.artmed.2011.11.007

the medical literature, where exemplar, noteworthy cases are often described.

The approach to elicitation developed here is limited to directed acyclic graphs (Dags), a tool commonly used to represent conditional independence relations [8] or to express causal relations among variables [9]. Conditional independence properties follow from causal relations [9] without further constraints as long as a causally adequate set of variables is considered and the whole causal structure is properly defined.

Since medical skills should permit a causal understanding of the patient problem, they are assumed to encompass the ability of specifying clinically relevant variables, like those representing the body's internal state, the manifestations and risk factors [10]. However, like other experts, medical practitioners can rarely abstract the variables that are directly related within the Dag, unless all the relevant variables are preliminarily offered [11].

The elicitation framework we proposed is built around an interview using automated questions which are adaptively formulated to build a Dag for a given case report by exploiting common patterns in medical causal reasoning. It can be also considered an instance of problem-based learning, where the starting point is typically the description of an individual clinical presentation [12], and which is meant to lead beginners and even expert physicians towards effective clinical reasoning on a patient's problem.

The interview is automated by means of four algorithms. Section 2.1 introduces the formal background deployed in the definition of the appropriate class of Dags for medical domains. Which relationships among medical variables are of interest is both described in medical terms and given a general graph representation in Section 2.2. Section 2.3 exploits such a graph to show that the algorithms are designed to submit all questions of clinical relevance. Two case studies selected from the medical literature illustrate actual steps and the output obtained by following the proposed approach. The final discussion encompasses possible applications of the elicitation framework and issues to be addressed in future research.

2. Methods

We take the perspective of eliciting medical doctor's beliefs about an individual patient. The description of the patient's problem is formatted as a clinical case report whose structure reflects the steps involved in the process of care, each one referring to a set of observations, judgments and decisions. For our analysis, only the clinical presentation is accessible, i.e. the observations initially collected. Reports typically provide details covering target variables of the clinical problem, such as the patient's main complaints, which may persist in the future, and any alarming conditions that put the patient's health at risk. General characteristics like sex and age are nearly always reported, and any associated chronic pathologies and medical treatments are noted.

Once a report is retrieved, the doctor is asked to provide the most plausible medical explanations of the evidence. This is tutored by an automated interview to gather the most relevant medical knowledge, which is expressed as a Dag. The elicitation framework we developed depends on the key assumption that doctors have enough medical background to represent current relevant events as variables of a causal Dag. Therefore, psychiatric problems are excluded, since difficulty in experimentation prevents a general agreement on psychiatric causal theories. Additionally, we restricted our attention to case reports concerning recent complaints due to either new acute disorders or novel manifestations of long-withstanding chronic conditions. This restriction avoids the explicit consideration of time and event history.

2.1. Notation and formal background on graphs

A graph $\mathcal{G} = (V, E)$ is a pair made by a finite set $V = \{v_1, v_2, \ldots, v_K\}$ of nodes and a collection of edges $E \subset V \times V$. The set E defines which nodes are linked by an edge, whether oriented or not: if $(v_i, v_j) \in E$ and $(v_j, v_i) \in E$ then the edge is undirected, $v_i - v_j$, otherwise if just one pair, say (v_i, v_j), is in the relation the edge is oriented, $v_i \to v_j$. The subset $pa(v_j) \subset V$ of nodes originating arrows reaching v_j is called parents set, $pa(v_j)$; the collection of children nodes $ch(v_j) \subset V$ is made by nodes reached by an arrow originated from v_i. A leaf node has at least one parent and no children. A root node has no parents and at least one child.

A subset of nodes in V is indicated by capital letters, like $V_A \subset V$ or $A \subset V$. A subset of edges $E_{A,B} \subset E$ refers to specific sets of parents and children, for example $E_{A,B}$ is a subset of edges where parents belong to V_A and children belong to V_B. A path is an ordered set of nodes $(v_0, \ldots, v_k)$ in which pairs v_i, v_{i+1} are connected by an edge. A directed path π is a path in which edges always meet head-to-tail. A directed path from a node $v_0 \in V_A$ to a node $v_k \in V_B$ is an element of the set $\pi(V_A, V_B)$ of all paths where $v_0 \in V_A$ and $v_k \in V_B$; the compact notation $v_i \in \pi(V_A, V_B)$ indicates that there is at least one path π in $\pi(V_A, V_B)$ that contains node v_i. The ancestors set $an(v_i)$ of node v_i contains nodes located on directed paths reaching v_i. The descendants set $de(v_i)$ of node v_i contains nodes reached by a directed path originated in node v_i. In a directed graph all edges are oriented. In a directed graph a cycle is a directed path in which $v_k = v_0$. A Dag is a directed graph without cycles.

In a connected Dag each node is origin or destination of at least one oriented edge. Useful collections of connected Dags may be often represented by a constraint Dag. A constraint Dag $\mathcal{C}$ for the class of Dags with nodes in V is a directed acyclic graph in which nodes are elements of a partition $\{c_1, c_2, \ldots\}$ of V. A constraint Dag is not trivial if at least one of its nodes contains two or more elements belonging to V (original nodes). An oriented edge $c_i \to c_j$ in $\mathcal{C}$ indicates that oriented edges $v_m \to v_n$ are allowed, where $v_m \in c_i$, $v_n \in c_j$. An oriented edge in a constraint Dag does not imply oriented edges for Dags on V, but overall the Dag on V has to be connected, thus an implicit constraint on absent edges is in force. All oriented edges not represented in $\mathcal{C}$ are forbidden and a Dag on V that contains forbidden edges is not contained in the set of Dags defined by $\mathcal{C}$. A constraint Dag $\mathcal{C}$ is called simple if it also forbids the presence of edges between nodes in V that are contained into the same root/leaf c_j of $\mathcal{C}$, thus if c_j is a leaf node in $\mathcal{C}$ than $v_i \in c_j$ are leaf nodes too.

A wider class of Dags is considered by changing the set of nodes V, for example if Dags without some nodes in V are also of interest. The reduction $\mathcal{C}_r$ of a constraint Dag $\mathcal{C}$ is a constraint Dag obtained by removing node c_j from $\mathcal{C}$. It follows that the original set V is also reduced by deleting nodes belonging to the removed $\mathcal{C}_j$, that is $v_i \in c_j$. Root nodes and leaf nodes are simply deleted from $\mathcal{C}$ with all their outgoing/incoming edges to obtain the reduced constraint Dag $\mathcal{C}_r$. If a node c_k with not empty $pa(c_k)$ and not empty $ch(c_k)$ is deleted than all incoming and outgoing edges are deleted but all directed edges in the set $\pi(pa(c_k), ch(c_k))$ are inserted in the reduced constraint Dag.

2.2. Dags consistent with medical background

The interview is designed to elicit medical variables MV and their plausible causal relations as long as they are viewed as relevant for the case on hand. A cardinal feature rests on questions based on standard medical terminology. The doctor introduces new variables by answering to questions regarding variables he has already included in the Dag. Elicited variables are indicated by the set V, wherein nodes follow a general classification based