Clinical Evidence Framework for Bayesian Networks

Barbaros Yet[1], Zane B. Perkins[2], Nigel R.M. Tai[3], and D. William R. Marsh[1]

[1]School of Electronic Engineering and Computer Science, Queen Mary University of London, UK

[2]Centre for Trauma Science, Queen Mary University of London, UK

[3]The Royal London Hospital, London, UK

Corresponding Author: Barbaros Yet

Contact Address: Department of Industrial Engineering, Hacettepe University, Ankara, Turkey

Email: barbaros.yet@hacettepe.edu.tr

Tel: +90 312 297 8950 /128

Abstract

Objectives

There is poor uptake of prognostic decision support models by clinicians regardless of their accuracy. There is evidence that this results from doubts about the basis of the model as the evidence behind clinical models is often not clear to anyone other than their developers. In this paper, we propose a framework for representing the evidence-base of a Bayesian network decision support model. The aim of this evidence framework is to be able to present all the clinical evidence alongside the BN itself. The evidence framework is capable of presenting supporting and conflicting evidence, and evidence associated with relevant but excluded factors. It also allows the completeness of the evidence to be queried. We illustrate this framework using a BN that has been previously developed to predict acute traumatic coagulopathy, a potentially fatal disorder of blood clotting, at early stages of trauma care.

Keywords: Bayesian Networks, Evidence Based Medicine, Prognostic Models, Clinical Decision Support, Knowledge Engineering

Graphical Abstract:
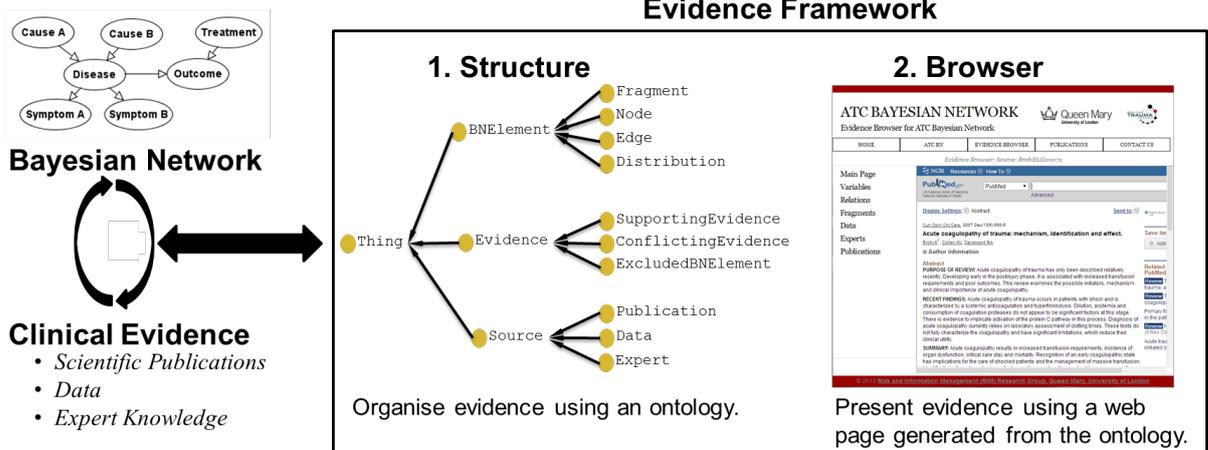
Highlights:

- We propose an evidence framework (EF) for clinical Bayesian network (BN) models.

- Our EF presents relevant clinical evidence alongside the BN itself.

- Evidence data is organised in a structured ontology, presented in a simple webpage.

- We illustrate the EF using a BN that has been previously developed for trauma care.

- Clinicians found the EF useful and practical in a qualitative evaluation session.

1 Introduction

Clinical decision support models are developed for a wide variety of tasks such as making recommendations, predictions or assessing cost-effectiveness. In this paper, our focus is on prognostic models that aim to predict a clinical outcome for individual patients using their clinical and background information (Moons et al., 2009b). Although a large number of prognostic models are developed and published, many are not adopted into clinical practice (Altman et al., 2009). Concerns about accuracy are considered to be one reason for this. Accuracy alone, however, does not ensure the use of a model (Jaspers et al., 2011; Moons et al., 2009a). While some models reporting accurate performance have not been adopted in clinical practice, others reporting mediocre performance have been widely used (Altman et al., 2009; Wyatt and Altman, 1995). Wyatt and Altman (1995) argue that useful prognostic models have 4 properties in common: clinical credibility, accuracy, generalisability, and ability to provide useful decision support. In order for a clinician to evaluate properties other than accuracy, the evidence supporting a model and the rationale behind its structure must be clear.

A Bayesian network (BN) is a powerful tool for building prognostic models based on clinical evidence (Lucas et al., 2004). The graphical structure of a BN is well suited for representing domain knowledge on causal and associational relations, and for combining evidence from publications and data (Fenton and Neil, 2012). Despite these advantages, a BN structure cannot present the corresponding clinical evidence in adequate detail. It shows only the names of the variables and directions of the relations between the variables. Moreover, variable names are often abbreviated and ambiguous, and the evidence relevant to the relations and variables are not shown at all. Consequently, clinical evidence supporting many BN models is clear only to their developers.

To improve the take up of BN models for clinical decision support, a method of organising and presenting clinical evidence is needed. This paper addresses this gap. We present an evidence framework that complements clinical BNs by representing the relevant clinical evidence and knowledge. The evidence framework does not aim to replace the representation of BNs; it is proposed as a complementary tool to clarify the clinical evidence relevant to different elements of a BN model. Our framework is composed of two parts: firstly, an ontology that organises the evidence, and

secondly, a browser that presents the evidence. We illustrate the evidence framework using a case study of a clinical BN that we have previously developed (Yet et al., 2014a). This BN accurately predicts a potentially fatal physiological disorder, called acute traumatic coagulopathy (ATC), in early stages of trauma care. The structure of the BN was built based on clinical knowledge and published evidence, and its parameters were defined using a patient dataset. More information on development and validation of the ATC BN is described by Yet et al., 2014a, and the clinical evidence relevant to the model can be browsed online on the model's website (ATCBN 2015). In this paper, we present the entire technical structure of the evidence framework and make a qualitative validation with its potential users.[1] The structure of the evidence framework, and the code for building the evidence browser are available in (Yet and Marsh 2015). In the remainder of this paper, Section 2 discusses the challenges of developing a clinical evidence framework for BNs. The components of our evidence framework are described in Section 3 (ontology) and Section 4 (browser). Section 5 presents a qualitative evaluation of the evidence framework. Sections 6 and 7 examine the related work and present our conclusions.

2   Challenges of Developing an Evidence Framework

In this section we summarise the challenges of organising and presenting clinical evidence relevant to a BN model.

- **Organising Evidence:** Clinical evidence corresponding to a BN must be recorded in a structured way to prevent ambiguities about its relevance, source and type. Below, we examine some issues that needs to be addressed in an evidence structure:

    1. The available evidence may be relevant to different parts of a BN: some may be relevant to a particular variable or relation whereas others may be relevant to an entire group of variables. A useful evidence framework would be able to accommodate all of these situations.

---

[1] A preliminary stage of the evidence framework has been briefly described in Yet et al., 2014b.

2. Different items of evidence may support different parts of a BN model. For example, a part of a BN may be based on data, and another part may be based on a scientific publication. The evidence framework must be able to accommodate and distinguish multiple items of evidence from different sources.

3. Conflicting evidence must not be ignored. In order to have a comprehensive evidence-base, all relevant evidence must be recorded even when it is relevant to excluded or conflicting factors.

4. Definitions of the variables and relations must be explicit.

5. The framework should allow the completeness of evidence to be queried. For example, it should be able to assist the user in identifying the parts of the BN with poor or no evidence.

- **Presenting Evidence:** Although the evidence framework may require a complex structure to organise the evidence, the users do not necessarily need to see the technical details of the structure when browsing the evidence. Evidence should be presented in a simple and user-friendly environment that is compatible with commonly available software such as web page browsers.

In order to overcome these challenges, we propose an evidence framework that is composed of two elements: 1) an evidence structure for organising the evidence data 2) a web page, which is automatically generated from the evidence structure, to present the evidence without showing the technical details of the underlying evidence structure. In the following sections, we describe the structure (Section 3) and browser (Section 4) parts of the evidence framework.

3 Evidence Structure

We use the web ontology language (OWL) (W3C, 2013) framework to model and organise the evidence data, and the Protégé software (version 4.3.0) (Knublauch et al., 2004) to create and populate the OWL ontologies. In Section 3.1, we give a brief introduction to ontologies. We present the structure of the evidence ontology in Section 3.2, illustrate how data is entered to this structure in Section 3.3, and show several queries to assess completeness of evidence in Section 3.4.

3.1 Overview of Ontologies

An ontology is a formal definition of the relations among terms (Berners-Lee et al., 2001), and OWL is a flexible language for building ontologies (W3C, 2013). Our primary reason for storing the evidence data in an ontology is their flexibility in building, modifying and querying a data structure. Our aim is to develop a general structure for organising and presenting clinical evidence for BNs but we do not assume that our evidence structure will satisfy the needs in all other BN applications. For example, some clinical applications may require additional types of evidence to be defined in the ontology structure. Ontologies offer a simple and robust framework for making such changes. Making similar modifications in relational databases is, however, difficult and time-consuming due to their highly structured schema and query system.

An OWL ontology is composed of individuals, classes, object properties and data properties. Individuals represent objects in the domain, and classes represent the sets that individuals belong. Properties define the relations between individuals: object properties define the relation between two objects, and data properties define the relation between an object and a data value. Throughout this paper, individuals are represented by diamonds, classes are represented by circles, object properties are represented squares, and data properties are represented by triangles as shown in Figure 1.

◆Individual
●Class
■Object Property
▲Data Property
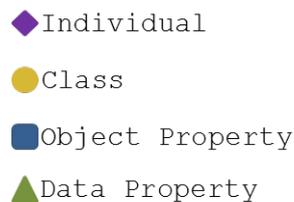
Figure 1 Representation of Ontology Elements

OWL contains many features for defining properties and objects. Members of a class and properties of an object can either be defined manually or be inferred from class and property characteristics. A thorough description of OWL is, however, beyond the scope of this paper, and the readers are referred to Allemang and Hendler (2011), and Segaran et al. (2009) for an introduction to ontology modelling and OWL. In the following section, we describe the classes, and properties of the evidence ontology.

3.2 Evidence Ontology

Our evidence structure is based on three main classes: BN element, evidence and source (see Figure 2). The BN element class contains the nodes, edges, fragments and probability distributions that form a BN. Each node in a BN represents a variable, and each edge represents a relation. Each node has a set of parameters that defines the conditional probability distribution between the node and its parents. We use the terms node and variable, edge and relation, and probability distribution and parameters interchangeably throughout this paper. The evidence class defines the type and statement of evidence. The source class describes the source of evidence; which can be a dataset, a domain expert or a scientific publication. Individuals of the evidence class may have multiple sources, for example, a relation in the BN structure may have evidence from several publications. In the remainder of this section, we describe the subclasses and properties related to the BN element, evidence and source classes.
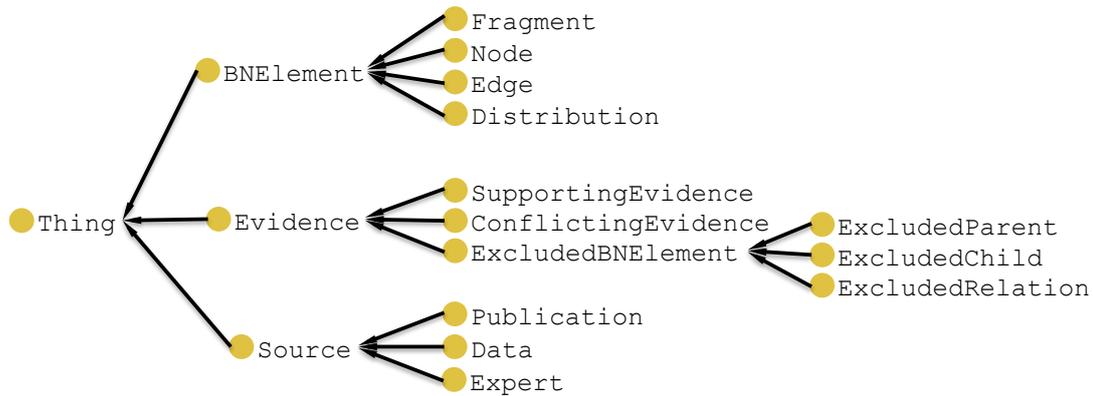
Figure 2 Class Hierarchy of the Evidence Ontology

3.2.1 BN element

The BN element class has 4 subclasses: fragment, node, edge and distribution. This section describes the object and data properties related to each of these subclasses.

Fragment

Several studies used BN substructures to assist BN development and to give a more concise summary of a BN model (Koller and Pfeffer, 1997; Laskey and Mahoney, 1997; Neil et al., 2000). A

BN substructure represents a part of the BN that describes an important concept in its domain. Our ontology defines BN substructures in a general way as 'BN fragments'. Any group of variables can be defined as a BN fragment in the evidence ontology. The complete structure of the BN is always defined as a BN fragment containing all other fragments and nodes in the BN.

The information about a BN fragment is stored within the fragment class. Figure 3 shows the object and data properties related to this class. A BN fragment may contain nodes and other – smaller – BN fragments. Members of a fragment can overlap with other fragments. An edge can also be a part of a BN fragment but it is not necessary to separately define the edges associated with a fragment as they can be inferred through the object properties of the node class (see Figure 5). Description of clinical knowledge modelled in the BN fragment, and evidence relevant to the entire BN fragment can be stored in the evidence ontology.
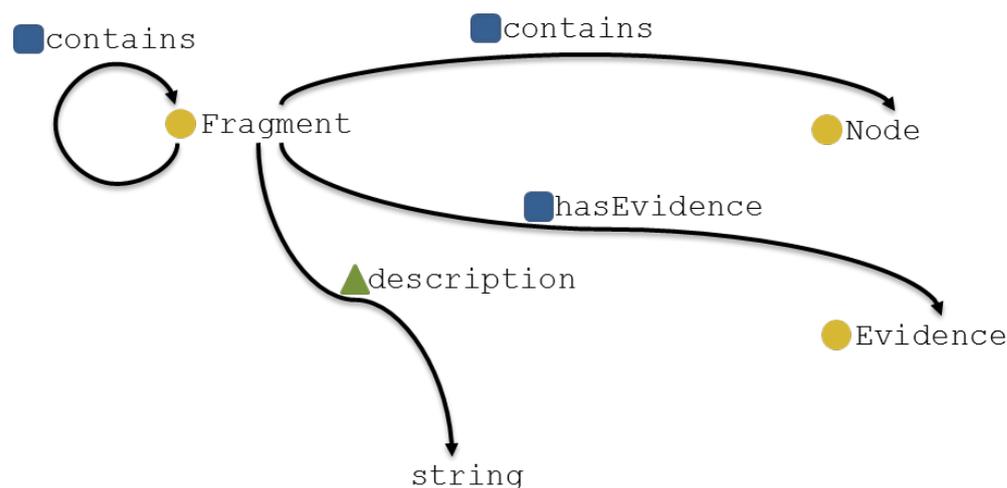
Figure 3 Object and Data Properties related to Fragment Class

Subclasses of the fragment class can be defined if certain type of BN fragments, such as idioms (Neil et al., 2000) and object oriented BNs (Koller and Pfeffer, 1997), are used repetitively in the BN. In the ATC BN, we use measurement idioms multiple times to define parts of the BN structure (Yet et al., 2014a). In order to distinguish the measurement idioms from other BN fragments, we could add a class called 'MeasurementIdiom' as a subclass of the fragment class (see Figure 4).

Figure 4 Subclasses of Fragment Class

Node

Individuals of the node class represent the variables in a BN. Knowledge modelled in a BN cannot be understood without having a clear understanding of the meaning of its variables. Variable names are often short, consisting of only a few letters, so they may not clearly describe the concept that is represented by the variable. For example, a BN variable named 'Heart Rate' may be sufficient to show that this variable represents a measurement of the patient's heart rate; but this name may not be descriptive enough if the time and location of this measurement is important for the use of the BN. In order to avoid such ambiguities, the evidence ontology stores a description of each BN variable (see Figure 5). Each variable has a conditional probability distribution that defines the relation between the variable and its parents. The states, parameters and the evidence for defining the probability distribution are defined within the associated distribution class.
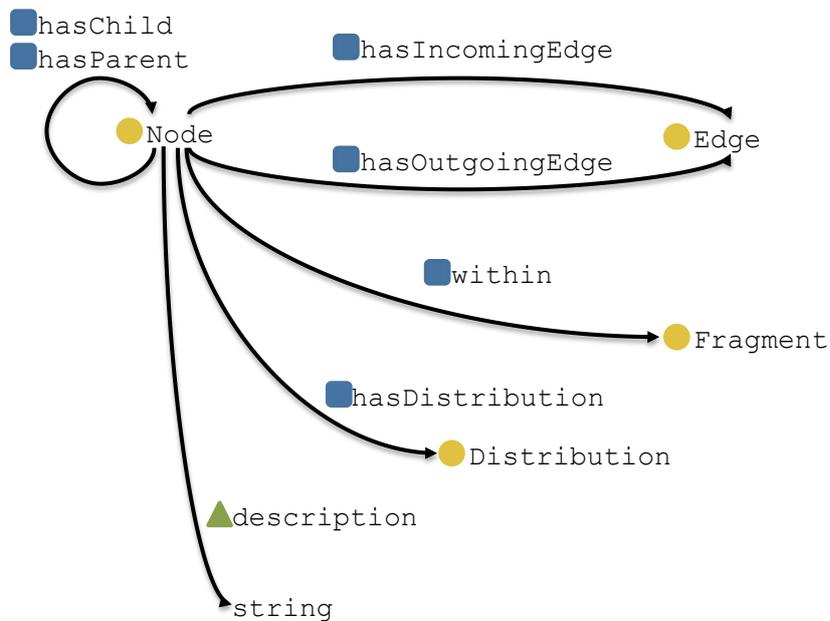

Figure 5 Object and Data Properties related to Node Class

Since the edges of a BN are directed, we have an object property for both the edges that are directed to the variable (hasIncomingEdge) and the edges that are directed away from the variable (hasOutgoingEdge) (see Figure 5). The hasParent and hasChild properties show the parent – child relations between the BN variables. These 4 properties about edge directions and parent – child relations are inferred from the object properties related to the edge class described in the remainder of this section. The within property shows the BN fragments that contain the variable.

Different sources can be used to define the parameters of different BN variables as BN parameters can be learnt locally. For example, we can use data to define the parameters of some variables, and use expert knowledge to elicit the parameters of others. For each variable, we record the source used for defining its parameters (see Figure 6).

Some variables may not be modelled in the BN even when evidence exists that they are relevant to the problem domain. For example, the knowledge engineers and domain experts may choose to exclude some variables to keep the BN simple. The evidence relevant to the excluded variables and relations is also recorded to have a comprehensive evidence base for the BN (see Figure 6).
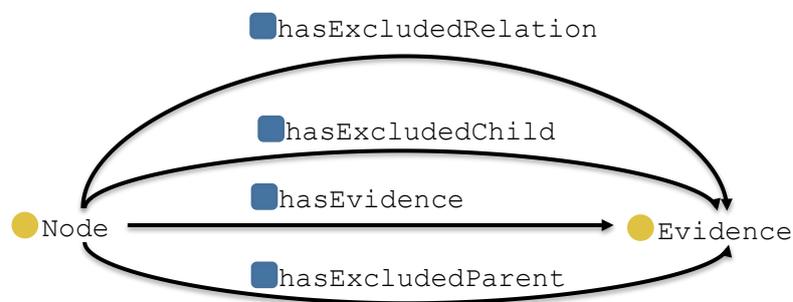

Figure 6 Object Properties related to Node, Evidence and Source Classes

Edge

BNs are directed acyclic graphs, therefore only one directed edge can exist between two variables. The variables that an edge connects are defined with two object properties in the evidence ontology: pointsTo shows the variable that the edge is directed to, and comesFrom shows the variable that the edge is directed away (see Figure 7).

The meaning of a modelled relation is shown by its supporting evidence. In order to include an edge in a BN model, the BN developer must have evidence that the relation exist between two variables. Therefore, all edges in an evidence-based BN model must have supporting evidence, which may come from scientific publications, data or domain experts. An edge does not require a separate description like a variable since evidence associated with the edge describes its meaning. When the BN structure is built with domain experts, evidence for the relations often comes from peer-reviewed scientific publications or expert knowledge. The justification for including the relation is recorded as the statement of evidence. In a purely data-driven BN, the evidence comes from the dataset only. It may be difficult to describe knowledge supporting the relations in data-driven models apart from saying that the variables were correlated in the data and showing the structure learning algorithm used (Daly et al. 2011).

The relevant score or independence test could also be shown when score or constraint based algorithms are used to learn the relation. The results of score-based algorithms can be shown together with the scores obtained from alternative learning algorithms, but the scores must be decomposable in order to calculate them for individual edges. Constraint-based methods can calculate the p-values for missing edges (Spirtes et al. 2000) and the p-value upper bounds (Li and Wang 2009; Tsamardinos and Brown 2008) for existing edges. Conflicting evidence may also exist for a relation in the BN. For example, one publication may claim that two variables are independent, whereas another publication may find correlation between those variables. If these variables are connected by an edge in the BN structure, in accordance with the latter publication, the conflicting evidence from the former publication must also be recorded in order to have a comprehensive evidence base.
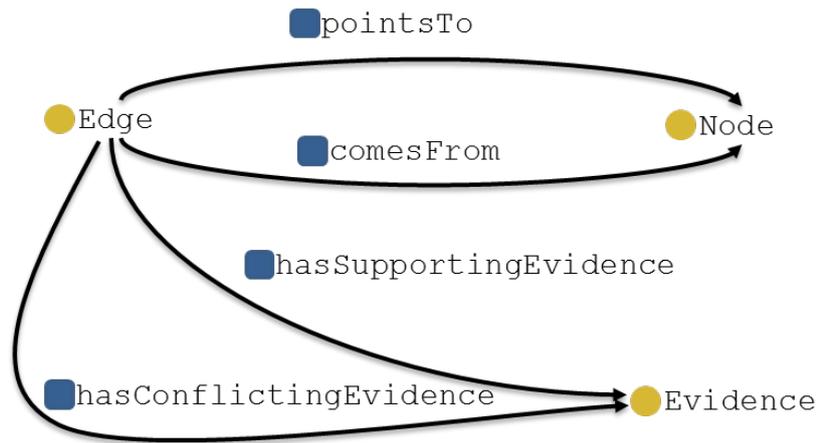
Figure 7 Object Properties related to Edge Class

An edge between two variables is present when those variables are directly dependent. In other words, no observation on other variables can make two variables conditionally independent when they are directly connected by an edge. However, there may be evidence relevant to two variables that are not directly linked in a BN model. For example, there may be evidence about correlation between two variables but not about a direct causal relation. In this case, the correlation may be modelled by a path between those variables in the BN, however representing the evidence of all paths in the BN could make the evidence base increasingly complex especially if the BN structure is large. A simple and effective approach to model indirect correlations is to use an 'excluded relation' evidence item in the evidence base (see Section 3.2.2 for more information on evidence items). For example, if there is evidence about correlation between A and B but in the BN model A and B is connected by a long path, we can introduce an 'excluded relation' evidence item between A and B to record the evidence about the correlation.

Inferred Properties

In an OWL ontology, class memberships and object properties can either be defined by the user or be inferred from other class and property characteristics. Although some inferred properties may be redundant, they can be useful when finding associations between different items of evidence or presenting the evidence in an online browser (see Section 4). For example, defining the direction of edges using the pointsTo and comesFrom properties is sufficient to define an edge but redundant properties such as hasParent and hasChild are useful when the BN is presented in the online browser.

13

The ontology can infer these properties from the pointsTo and comesFrom properties. In order to define the directed relation A → B between two variables A and B, we define an individual for the edge class, e.g. Edge1, and we define that Edge1 pointsTo B and comesFrom A. In our ontology, the inverse of the comesFrom property is the hasOutgoingEdge property. By using this, the ontology can infer that A hasOutgoingEdge Edge1. The hasChild property can be inferred from the combination of the hasOutgoingEdge and pointsTo properties. Since the variable A hasOutgoingEdge Edge 1 and Edge1 pointsTo the variable B, the ontology can infer that A hasChild B. Similarly, the ontology can infer the hasParent property from the combination of the comesFrom and hasIncomingEdge.

Distribution

Each variable in a BN has a probability distribution that defines the relation between the variable and its parents. The probability distribution of a variable can be discrete or continuous and it can be defined in different formats. If a variable and its parents are discrete, a node probability table (NPT) is commonly used to define their conditional probability distribution. If the child variable is continuous, using a mathematical expression is a convenient way of defining its probability distribution. Mixture distributions of a continuous variable conditioned on discrete variables can be defined with an array that contains an expression for every state combination of its parents. We call these arrays partitioned expressions. Other formats for defining probability distributions can also be added as subclasses of the distribution class.

The method and resource used to define a probability distribution is recorded as supporting evidence for the distribution. A probability distribution can also have conflicting evidence: for example, the views of different experts that have conflicting views on the probability distribution of a variable can be recorded as conflicting evidence in the ontology.
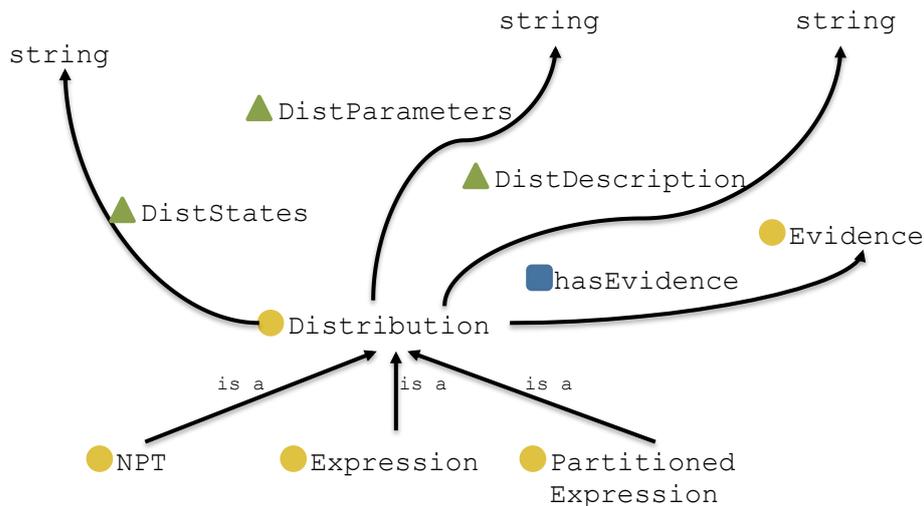
3.2.2 Evidence

Figure 8 shows the properties related to the evidence class. The evidence class organises the data about the type and statement of evidence relevant to the BN elements. There are 3 types of evidence in our ontology:

1. **Supporting Evidence:** This class contains the statement of evidence that supports the elements of a BN model.

2. **Conflicting Evidence:** Evidence conflicting with the BN model is recorded in this class. For example, two variables that are found to be independent in a scientific study may not be modelled as independent in the BN model. In this case, the results of the scientific study must be recorded as conflicting evidence in the evidence ontology. Conflicting expert opinions may also be recorded in this category.

3. **Excluded BN Element:** Evidence about relevant variables or relations that are not included in the BN model is recorded within this class. The statement of evidence relevant to an excluded BN element and the justification for not including the element should be recorded. The excluded BN element class has 3 subclasses related to the type of the BN element excluded. The excluded child and parent classes indicate that a variable in the BN may have an effect and cause that is not modelled in the BN. The excluded relation class indicates that a direct relation may exist between two variables that are not directly linked in the BN.

Each type of evidence is a subclass of the evidence class. The subclasses can be inferred from the type of the object property that refers to the evidence. For example, any evidence recorded by the hasSupportingEvidence property is inferred as a member of the SupportingEvidence class. Similarly, any evidence that is recorded by the hasConflictingEvidence or hasExcludedBNElement properties are inferred as the members of the ConflictingEvidence or ExludedBNElement classes respectively.
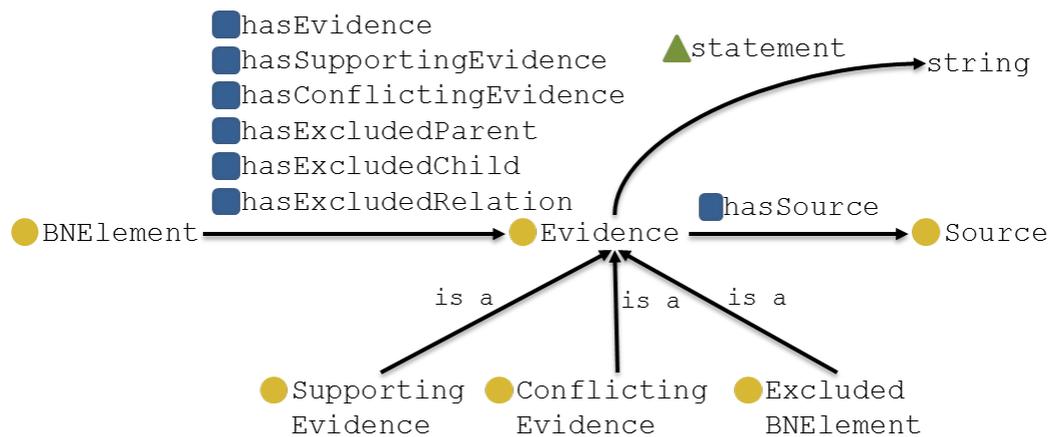

Figure 8 Object and Data Properties related to Evidence Class

A BN element can have multiple items of evidence, and an item of evidence can have multiple sources. For example, two publications that are stating similar results about a relation in the BN can be recorded as a single evidence item with two sources (publications). Two publications that are discussing different aspects of a relation can be recorded as two separate items of evidence with one source each. More examples about recording evidence and source are shown in Section 3.3. The following section describes the ontology class for the source of evidence.

3.2.3 Source

The source class contains information about the publication, expert opinion or data providing evidence. The subclasses, object and data properties of the source class are shown in Figure 9. The source class has three subclasses as shown below:

1. **Publication:** The members of this class are scientific publications. Information for referring these publications, such as digital object identifier (DOI) and PubMed identification number

(PMID), can be stored in this class. A PMID is a unique identifier assigned that makes it convenient to find the publications in the PubMed database.

2. **Expert:** The domain experts' credentials and contact information can be recorded in this class.

3. **Data:** Evidence supporting or conflicting with the model may come from data. The individuals of this class contain information about the details of the dataset including the sample size and the method of collecting the data.
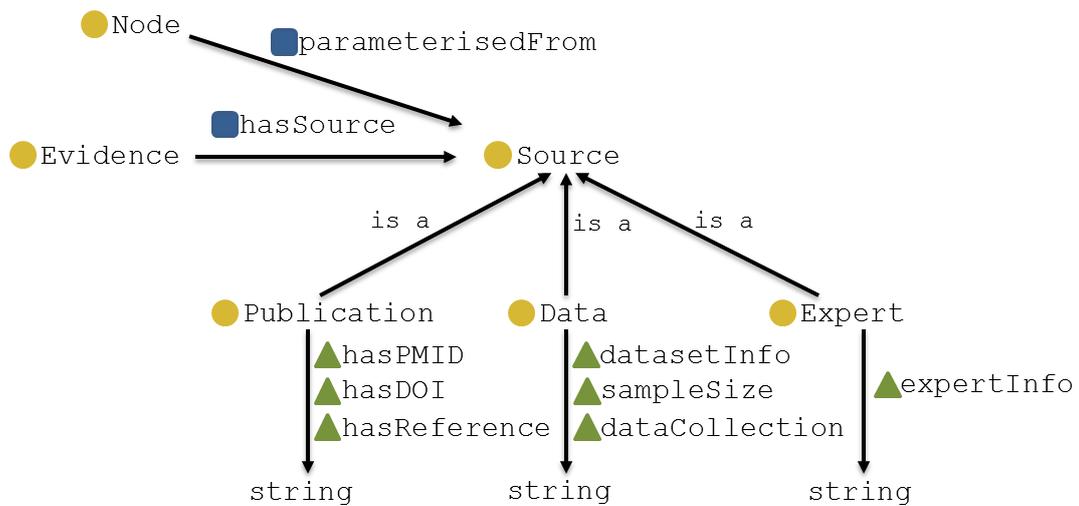

Figure 9 Object and Data Properties related to Source Class

3.3 Instantiation of the Evidence Ontology

In this section, we illustrate an instantiation of the evidence ontology by using a simplified version of the ATC BN that has 4 variables and 3 edges (see Figure 10).
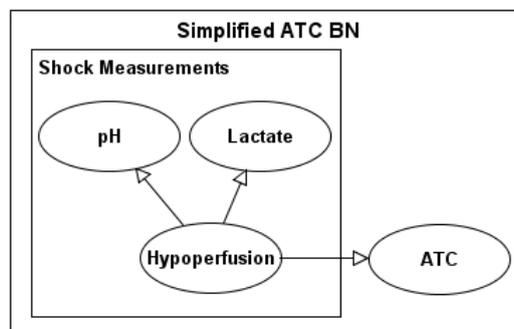
Figure 10 Simplified ATC BN

3.3.1 *Fragment*

The simplified ATC BN in Figure 10 contains a BN fragment called shock measurements. The state of the hypoperfusion variable cannot be directly observed in clinical practice; it is estimated by several tests and observations including the lactate and pH levels in blood. The shock measurements fragment models the relation between hypoperfusion and its measurements.

Table 1 Defining Shock Fragment

◆ ShockMeasurements	

Type:
🟡 BNElement
  🟡 Fragment
    🟡 MeasurementIdiom
	Object Properties:
	🟦 within SimplifiedATCBN
	🟦 contains Hypoperfusion
	🟦 contains Lactate
	🟦 contains pH
	
	Data Properties:
	🔺 description "This part of the model estimates…"

Table 1 shows the classes, object and data properties related to the shock measurements fragment. The structure of the shock measurements fragment is modelled using a pre-defined BN structure called measurement idiom (Neil et al., 2000). We use the object property 'contains' to show the variables and fragments within this fragment, and 'within' to show the larger fragments that contain this fragment. A free-text description of the fragment is recorded using a data property. The entire BN is also defined as a BN fragment (SimplifiedATCBN) that contains all other variables and fragments in the evidence ontology (see Table 2). The aims and scope of the model is recorded using the 'description' data property. The search strategy used for finding the relevant publications is recorded as a supporting evidence item for the entire BN model.

Table 2 Entire BN model defined as a member of the fragment class

◆ SimplifiedATCBN	
Type: 🟡 BNElement 　🟡 Fragment	Object Properties: 🟦 contains SimplifiedATCBN 🟦 contains Hypoperfusion 🟦 contains Lactate 🟦 contains pH 🟦 contains ATC 🟦 hasSupportingEvidence SearchStrategy Data Properties: 🔺 description "The primary aim of the ATC BN is to predict…"

3.3.2 Variable

Table 3 shows the data related to the Hypoperfusion variable in the evidence ontology. This variable is within the Shock fragment, and it has three children: ATC, Lactate and pH. The probability distribution of this variable is defined in the 'HypoNPT' individual. The description of the variable is recorded by the 'description' data property. The edges between the Hypoperfusion variable and its children are modelled by the 'hasOutgoingEdge' object property that shows the edges that is directed away from the Hypoperfusion variable. For example, 'HypoToLactate' is a member of the edge class that represents the edge Hypoperfusion→Lactate. The evidence supporting the relation between Hypoperfusion and Lactate variables is defined under the edge individual.

Table 3 Defining the Hypoperfusion Variable

◆ Hypoperfusion	
Type: 🟡 BNElement 　🟡 Variable	Object Properties: 🟦 within Shock 🟦 hasChild ATC 🟦 hasChild Lactate 🟦 hasChild pH 🟦 hasOutgoingEdge HypoToATC 🟦 hasOutgoingEdge HypoToLactate 🟦 hasOutgoingEdge HypoToPH 🟦 hasDistribution HypoNPT Data Properties: 🔺 description "The degree of inadequate oxygen…"

Table 4 Defining the Lactate Variable

◆ Lactate	
Type: 🟡 `BNElement` 🟡 `Variable`	**Object Properties:** 🟦 `within Shock` 🟦 `hasParent Hypoperfusion` 🟦 `hasIncomingEdge HypoToLactate` 🟦 `hasExcludedParent LactateExcParEv` 🟦 `hasDistribution LactatePartExpr` **Data Properties:** 🔺 `description "The amount of lactate in…"`

Table 4 shows the information recorded about the lactate variable. The probability distribution of lactate is defined in the 'LactatePartExpr' distribution (see Section 3.3.3). There is evidence that some factors affecting lactate are not included in this model. This information is recorded in an item of evidence named 'LactateExParEv1' belonging to the evidence class (see Section 3.3.4).

3.3.3 Distribution

The probability distribution of the hypoperfusion variable is shown in Table 5. The dataset and method used for defining this distribution is recorded as supporting evidence for this distribution. The states of the Hypoperfusion variable is recorded by the 'DistStates' data property, and the probability distribution of these states are shown in the same order as the states by using the 'DistParameters' data property. In this example, we use the BIF format for recording the NPTs in the ontology but other formats could also be used (Cozman, 1998). The 'DistDescription' data property shows that Hypoperfusion is not conditioned on any variables in the simplified ATC BN.

Table 5 Defining the NPT of Hypoperfusion

◆ HypoNPT	
Type: 🟡 `BNElement` 🟡 `Distribution` 🟡 `NPT`	**Object Properties:** 🟦 `hasSupportingEvidence ParamEv`

Data Properties:
▲ DistDescription "P(Hypoperfusion)"
▲ DistStates "None,Compensated, Uncompensated"
▲ DistParameters "0.74 0.17 0.09"

Table 6 shows the probability distribution of the Lactate variable. Lactate is a continuous variable conditioned on Hypoperfusuion as shown by the 'DistStates' and 'DistDescription' data properties respectively. Lactate has a normal mixture distribution, and the parameters of this distribution are recorded by the 'DistParameters' property.

Table 6 Defining the Partitioned Expression of Lactate

◆ LactatePartExpr	
Type: 🟡 BNElement 　🟡 Distribution 　　🟡 Partitioned Expression	Object Properties: 🟦 hasSupportingEvidence ParamEv Data Properties: ▲ DistDescription "P(Lactate|Hypoperfusion)" ▲ DistStates "Continuous" ▲ DistParameters " (None) Normal(1.7,3); (Compensated) Normal(4.5,1.8); (Uncompensated) Normal(8.2,23.2);"

3.3.4 Evidence

Two publications indicated that lactate is an important marker of the degree of hypoperfusion. We recorded this as an item of evidence supporting the edge Hypoperfusion → Lactate in the ATC BN (see Table 7). The details of the publications are described under the source class in Section 3.3.5.

Table 7 Supporting Evidence of Lactate Variable

◆ LactateSuppEv	
Type: 🟡 Evidence 　🟡 SupportingEvidence	Object Properties: 🟦 hasSource RixenEtAl2005 🟦 hasSource VanDrommeEtAl2010

Data Properties:
▲ statement "Lactate is produced during anaerobic…"

The parameters of the conditional probability distribution of the lactate variable were learned from the dataset using the Expectation-Maximisation (EM) algorithm. This was noted as a supporting evidence for the 'LactatePartExpr' distribution object (see Table 8).

Table 8 Supporting Evidence of Lactate Parameters

◆ ParamEv	
Type: 🟡 Evidence 　🟡 SupportingEvidence	Object Properties: 🟦 hasSource RLHDataset Data Properties: ▲ statement "Parameters of this variable was learned from the RLH dataset using the EM algorithm…"

One publication indicated that several factors, including excessive alcohol use, can affect lactate levels independent from the effects of hypoperfusion but these factors were not included in the ATC BN. We recorded this as evidence for an excluded parent variable for Hypoperfusion (see Table 9).

Table 9 Evidence about Excluded Parent

◆ LactateExcParEv	
Type: 🟡 Evidence 　🟡 ExcludedBNElement 　　🟡 ExcludedParent	Object Properties: 🟦 hasSource RixenEtAl2005 Data Properties: ▲ statement "…excessive alcohol use can increase the lactate levels …"

Table 10 Search Strategy for Finding the Relevant Publications

◆ SearchStrategy	
Type: 🟡 Evidence 🟡 SupportingEvidence	**Object Properties:** 🟦 hasSource ZanePerkins **Data Properties:** 🔺 statement "… relevant publications were identified by an electronic search of the EMBASE and PUBMED database using combinations of the terms……"

The search strategy, which was used to identify the relevant publications from the clinical literature, is recorded as a supporting evidence item associated with the entire BN model (see Table 10). This item describes the methodology of the search such as the search keywords and exclusion criteria that were used. The clinicians who conducted the literature search are recorded as the source of this evidence item.

3.3.5 Source

Table 11 shows the details of the publication related to the items of evidence in Table 7 and Table 9. The publication is recorded as a member of the source class with its referencing details and PMID.

Table 11 Publication Source

◆ RixenEtA12005	
Type: 🟡 Source 🟡 Publication	**Data Properties:** 🔺 hasPMID 16277731 🔺 hasRefDetails "D. Rixen and J. H. Siegel, Critical Care, vol. 9, no. 5, p. 441, 2005 …"

Table 12 shows the details of the database related to the item of evidence in Table 8. We recorded the description, sample size and method of collecting the dataset.

Table 12 Data Source

◆ RLHDataset	
Type: 🟡 Source 🟡 Data	Data Properties: 🔺 datasetInfo "A dataset of 600 patients who were treated at the…" 🔺 dataCollection "Prospective Observational" 🔺 sampleSize 600

3.4 Completeness Queries using SPARQL Query Language

The parts of a clinical BN that lack evidence may be interesting to clinicians who review the BN. Therefore, an efficient way of assessing the completeness of evidence can be a useful feature for the evidence framework. SPARQL is a query language that can retrieve the data stored in an OWL ontology. In this section, we show several SPARQL queries to find the BN elements that have or lack evidence. It is beyond the focus of this paper to give a comprehensive description of SPARQL, a thorough introduction is given by Allemang & Hendler (2010).

The simplest form of query in SPARQL is the *SELECT* query, which extracts the data from the ontology and presents them in a table format. The *SELECT* query is followed by the *WHERE* block that limits the query by a question pattern. The *DISTINCT* keyword to filters out the duplicate results from the edges that have multiple items of evidence. For example, we can get the list of the edges that have supporting evidence by using the following query:

```
SELECT DISTINCT ?x

   WHERE { ?x a :Edge.

        ?x :hasSupportingEvidence ?evidence. }
```

The BN elements without evidence can be retrieved by using the *MINUS* keyword within the *WHERE* construct. The *UNSAID* keyword can be used as an alternative to *MINUS* keyword in SPARQL 1.1. The following SPARQL query can be used to find the edges that do not have supporting evidence.

```
SELECT DISTINCT ?x

  WHERE { ?x a :Edge.

        MINUS { ?x :hasSupportingEvidence ?evidence . }}
```

The BN elements that are not covered by any fragment can also be retrieved by using the *MINUS* keyword. The following SPARQL query shows the nodes that are not covered by any fragment.

```
SELECT DISTINCT ?x

  WHERE { ?x a :Node.

        MINUS { ?fragment :contains ?x . }}
```

4 Browsing Evidence

The evidence ontology is well suited for organising and querying evidence but it is not a convenient tool for browsing evidence especially when the user is not proficient with the ontology language. Therefore, our ontology framework automatically prepares a web page (HTML files) for browsing evidence after the data is entered to the evidence ontology (a Protégé OWL file).

The web page generator is not specific to a particular BN; it can generate a web page for any BN model given that evidence is entered to Protégé OWL using the structure presented in Section 3.2. In the remainder of this section, we use the complete version of the ATC BN (ATC BN 2015; Yet et al., 2014a) as a case study to demonstrate a web page generated from the evidence ontology. We call this web page the evidence browser for the ATC BN (ATCBN 2015). An instantiation of the evidence browser for the Asia BN is also available in (Yet and Marsh 2015).

Figure 11 ATC BN fragment shown in the Evidence Browser

The ATC BN has 7 fragments in total (see Figure 11):

- The entire BN structure ('ATCBN')

- The mortality, injury, coagulopathy and shock fragments.

- Two measurement idioms that models the relation between the latent variables and their measurements ('ShockMeasurements' and 'ATCMeasurements'). These measurement idiom fragments exist within the shock and coagulopathy fragments.

Figure 11 is the presentation of the ATC BN fragment – the entire BN structure – in the evidence browser. The link to the development methodology of the model and the description of the search strategy used for identifying relevant published evidence is shown at this page. This page also contains the links to the fragments and variables within the ATC BN. Figure 12 shows the presentation of the Hypoperfusion variable in the evidence browser. The browser shows the

description of the variable, its states and relevant evidence. In this example, there is evidence of excluded child variables of this variable. The 'Hypoperfusion ↛ Respiratory Rate' item shows that the degree of Hypoperfusion can be also be measured with other markers that are not included in the ATC BN. These markers could be added as a child of the hypoperfusion variable. The reason for not including these variables is also described.

Figure 12 Hypoperfusion variable shown in the evidence browser

A link to the conditional probability distribution (HypoNPT) of the hypoperfusion variable is also shown in Figure 12. When the user clicks to this link, the parameters of this conditional probability distribution are shown (see Figure 13). The evidence supporting the definition of the states and parameters of these variables is also shown in this page.

The variables that are directly related to the hypoperfusion variable, as its child or parent, are also listed in Figure 12. Each variable in this list is linked to a page showing the evidence about the

relation. For example, when the user clicks to the Hypoperfusion → ATC link in the 'relations' section in Figure 12, the evidence related to this relation is shown (see Figure 14).

Evidence Browser: Distribution: HypoNPT

Main Page	
Variables	
Relations	
Fragments	
Data	
Experts	
Publications	

Probability Distribution:	**P(Hypoperfusion|BBP)**
Distribution Type:	Node Probability Table
States:	None, Compensated, Uncompensated
Evidence:	• Severity of hypoperfusion may be categorised as compensated / uncompensated (Williams et al., 2013). The method for classifying the hypoperfusion states is described by Yet et al. (2014). **[References]** WilliamsEtAl2013 YetEtAl2014 • Parameters of this variable has been learned from the Royal London Hospital's Trauma Dataset using the Expectation-Maximisation algortihm. **[References]** RLHDataset

P(Hypoperfusion|BBP)	None	Compensated	Uncompensated
(0)	0.89	0.08	0.03
(I)	0.66	0.23	0.11
(II)	0.37	0.27	0.36
(III)	0.03	0.29	0.68
(IV)	0.01	0.04	0.95

Figure 13 Conditional Probability Distribution of Hypoperfusion shown in Evidence-Browser

There are 4 items of evidence relevant to the Hypoperfusion → ATC relation (see Figure 14). The user can click any of the references in Figure 14 to examine the source of the evidence statements. For example, if the user clicks to 'BrohiEtAl2007a', the browser shows the details of this publication by connecting to the PubMed database (see Figure 15). The browser uses the PMID stored in the evidence ontology to find the publication in the PubMed database. If the source is a domain expert or a dataset, the browser shows the related information such as the institution and credentials of the expert or the description of the dataset.

Figure 14 Relation between Hypoperfusion and ATC shown in the Evidence Browser

Figure 15 A Referred Publication shown in the Evidence Browser

5 Qualitative Evaluation

We conducted a qualitative evaluation session to examine the opinions of clinicians about the practical use and potential contributions of the evidence framework. The aims of this session were to examine 1) a clinician's confidence in using models with clear and understandable evidence-base 2) the practical usefulness of the evidence framework. In the remainder of this section we describe the methodology and results of the qualitative evaluation session.

5.1 Methodology

The evaluation session took place in the Royal London Hospital, UK. Since the ATC BN is designed to be used at early stages of trauma care, the potential users of the model have different clinical specialties. The evaluation session was purposely conducted at a regular trauma meeting where clinicians from a wide variety of specialties attend. The clinicians, who were involved in development of the ATC BN, did not participate in this session. The total duration of the trauma meeting was 1 hour. The session was composed of two parts, each part lasted 30 minutes. In the first part, we presented a tutorial about BN models, and illustrated the main differences between an evidence-based BN model and other commonly used modelling approaches such as regression models. This was followed by a demonstration of the evidence browser using the ATC BN as a case study. We showed the items of evidence supporting different parts of the ATC BN using the online browser. The evidence framework was available online during the session, and the website address was given to the participants in case they would like to test the use of the framework. One-to-one tutorials were not given due to time limitation. In the second part, we collected qualitative data by using a questionnaire composed of 18 multiple-choice questions. Each question had a statement and a 5 level agreement scale, ranging from 1 "strongly disagree" to 5 "strongly agree". Participants marked the scale representing their degree of agreement with the corresponding statement. Figure 16 shows an example question. There was also one open-ended question for participants to comment and give suggestions for improving the evidence framework. We stayed with the participants to clarify any unclear points while they answered the questionnaire. Participants were asked to specify their clinical

specialisation and invited to leave their email addresses in case they were interested in participating further evaluations.

Questionnaire Item	Scale				
	Strongly Disagree	Disagree	Neutral	Agree	Strongly Agree
The evidence framework could be used in practice.	1	2	3	4	5

Figure 16 A sample statement from the questionnaire

5.2 Results

Sixteen participants attended the evaluation session, 15 of 16 (94%) participants answered all 18 questions and 1 (6%) participant answered 13 of 18 (72%) questions. Of the participants, 15 (94%) participants were clinicians and 1 (6%) was an occupational therapist doing research on trauma models and outcomes. Among the participating clinicians, 5 (33%) were emergency medicine specialists, 5 (33%) were anaesthesia specialists, 2 (13%) were general surgeons, 1 (6%) was a trauma surgeon, and 1 (6%) was an orthopaedic surgeon. These specialties are considered to be the primary users of the ATC BN.

Both clarity of supporting evidence and published predictive accuracy was considered to be essential factors for practical use of a prognostic model. All participants indicated that they prefer models that make sense clinically and are compatible with the best available evidence. Similarly, all participants indicated that they are more likely to use models that have high predictive performance. A way to review the evidence-base of a model is also seen as an attractive feature. All participants agree that a way to review the evidence-base of a prognostic model is needed, and over 80% of them think that a way to suggest modifications following new evidence would be useful. Figure 16 shows the mean and standard deviation of the response for these questions as black points and errors bars respectively.
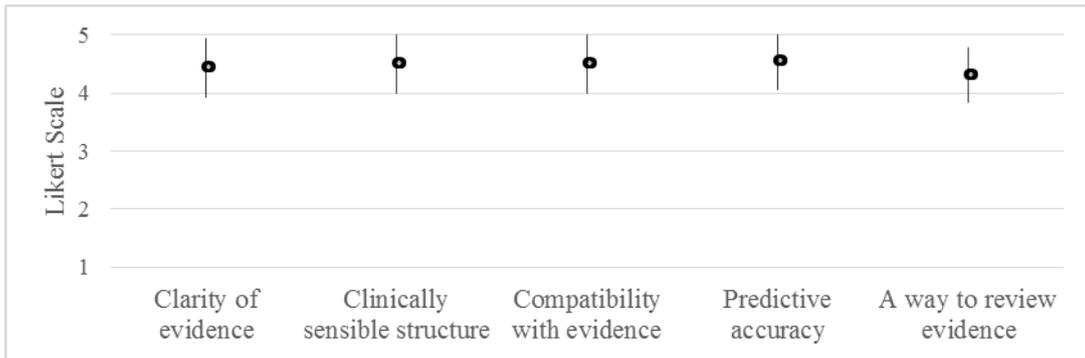
Figure 17. Questionnaire response relevant to individual importance of different factors

Since the questions about performance and clarity of evidence may be leading when asked individually, we assessed the relative importance of these factors to each other. We asked whether accuracy would be an important factor for clinical use when evidence-base is not clear or, in other words, the model is developed by a 'black-box' approach. In this case, the majority of the participants (75%) indicated that they are not likely to use such models in practice even when its published performance is good. Moreover, 85% of the participants indicated that they are not likely to models with high predictive performance if they contain variables that are not aligned with the best available clinical evidence. In other words, predictive accuracy alone is not enough for the majority of the clinicians to use a model especially if its evidence-base is not clear (see Figure 17).
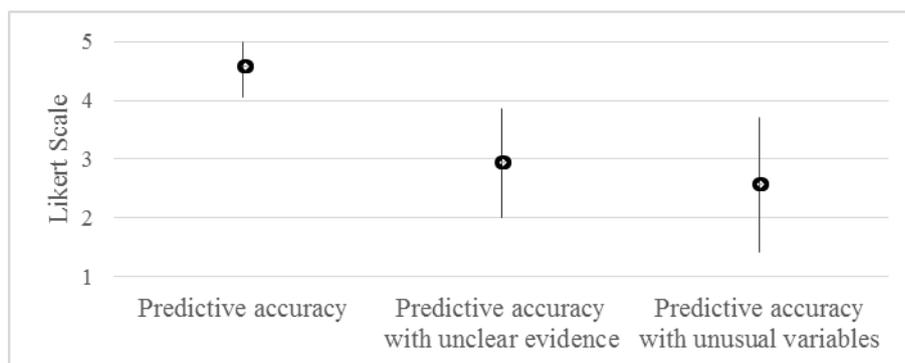
Figure 18 Questionnaire response relevant to individual and relative importance of predictive accuracy

Finally, the evidence framework was considered to be a useful and practical tool for clinical models. Over 85% of the participants indicated that our evidence framework improves the current practice in prognostic models, and over 90% thinks that this tool could be used in practice.

6 Related Work

Several other frameworks have been proposed to describe knowledge behind a BN model. Antal et al. (2001) initially propose annotated BNs as a tool to describe background knowledge for each node, edge, state and fragment of BNs. However, the structure of annotated BNs is not described in detail by Antal et al. (2001). In a following paper, Antal et al. (2004) simplify annotated BNs and define it as "a directed, acyclic graph whose nodes are annotated with local probabilistic models (as in standard Bayesian networks) and with textual attributes". In other words, they use annotations for only variables. They primarily use annotated BNs for automated learning of BNs from textual information in medical literature (Antal et al., 2004). Helsper and van Der Gaag (2007) use ontologies to clarify the elicited knowledge used for building a BN. They use a semi-formal ontology language that includes tables, figures and natural language descriptions. Their ontology documents expert knowledge from static, dynamic and definitional perspectives. Causal and temporal relations are described using Rieger and Grinberg's representation (1977). Helsper and van der Gaag (2002) also propose a methodology to derive a BN structure from this ontology however the method was not used to develop the model in their case-study since the ontology is prepared retrospectively from an existing BN. Van der Gaag and Tabachneck-Schijf (2010) extend Helsper and van Der Gaag's ontology framework (2007) for describing BNs that are built for similar tasks.

Both Helsper and van Der Gaag's ontologies (2007) and Antal et al.'s annotated BNs (2001) aim to clarify knowledge behind BNs but these frameworks do not aim to show the link between the model and evidence. Helsper and van Der Gaag (2007) focus on describing knowledge elicited from experts in more detail. For example, their ontology has numerous types of causal and temporal relations that can distinguish between continuous and one-shot causal events. However, Helsper and van der Gaag do not show the link between the model and evidence. Knowledge for their ontology is elicited from experts. Evidence from publications or conflicting expert opinions is not shown. Although the aims of the initial version of annotated BNs (Antal et al., 2001) are similar to our evidence framework, it lacks the structure to store information about the type and source of evidence, and excluded factors. These features are crucial for clarity and completeness of evidence. Moreover

the structure of the initial version is not described thoroughly in the paper. The final version of annotated BNs (Antal et al., 2004) can only store textual information relevant to variables as it is primarily designed to learn BN structure from textual information. The information relevant to edges or fragments is not supported. Several studies have developed methodologies to use existing ontologies for automated construction of BNs. Devitt et al. (2006) propose a methodology for automated construction of a BN from ontologies in the telecommunication networks domain. In this method, an ontology specific to the BN is derived from a more general ontology, and the BN is automatically generated from the specific ontology. Similarly, Sadeghi et al. (2005) build an ontology that is more specific to the problem domain using the concepts from the unified medical language system (UMLS) (Bodenreider, 2004) and they learn a BN based on this specific ontology. UMLS is a complex medical terminology system but it lacks information about the causal relations between clinical factors. Ishak et al. (2011) presents a set of rules that transform an ontology into a preliminary OOBN structure. Bucci et al. (2011) uses ontologies to build BNs on a predefined hierarchical structure in medical diagnosis domain. Fenz (2012) uses a semi-automated methodology to generate BNs from ontologies. In their methodology, the experts review an ontology and identify the nodes and states of a BN using the classes and individuals in the ontology. Afterwards they identify the edges in the BN using the object properties in the ontology. Fenz proposes a technique that uses the weights defined in the ontology to parameterise conditional probability tables in a similar approach to the parameterisation of ranked nodes (Fenton et al., 2007). An evidence-based BN requires justification for the relations modelled in the BN. Each edge in the BN should be supported with clinical studies, expert opinion or data. Many clinical ontologies are defined as clinical terminologies therefore they do not contain detailed information about causal and associational relations, and about medical publications and datasets relevant to those relations. Therefore, the reviewed studies about automated BN construction from ontologies are not aligned with the aims of our evidence framework.

Another active field of research that combines probabilistic models and ontologies focuses on extending ontologies with BNs and similar probabilistic models to represent uncertain knowledge in an ontology. Studies in this field use probabilistic models as a supporting reasoning mechanism for

the ontology. Ding and Peng (2004) propose additional mark-ups to OWL to represent probabilistic information. They present a set of rules that transforms an ontology into a BN. Similarly, Yang and Calmet (2005) present a BN extension to OWL that is also able to cope with multinomial variables. Costa et al. (2008) extend OWL to express uncertainty using multi entity BNs (Laskey and Costa, 2005) which is a combination of first-order logic and Bayesian reasoning. Kuo et al. (2013) builds BNs on ontologies to reason with undefined properties. Richardson and Domingos (2006) combines first-order logic and probabilistic models in Markov Logic Networks to represent uncertain knowledge. Zheng et al. (2008) propose mark-ups to transform an ontology representing clinical concepts into a BN. The aims of these studies are different from the aims of our evidence framework. These studies extend ontology languages, such as OWL, so that the ontologies are able to cope with uncertainties about class membership, object properties and their other features. The evidence framework uses the ontology as a complementary tool to clarify knowledge and evidence behind a decision support model, and it does not require the ontology to deal with uncertainty.

7 Conclusion

This paper proposed an evidence framework that complements clinical BNs by representing relevant clinical evidence and knowledge. The proposed framework is composed of two parts: an ontology that organises evidence relevant to different elements of a BN, and a browser that presents the BN and evidence to clinicians. To illustrate the evidence framework, we used a BN that was previously developed to predict a potentially fatal physiological disorder called ATC. We also conducted a qualitative validation session with 16 clinical consultants, registrars and researchers to examine the practical usefulness of evidence-based BN models and our evidence framework.

Our evidence framework is able to organise and present clinical evidence with more detail than it is possible with the existing BN representation or previously proposed annotation techniques. The evidence framework can store various types of evidence including evidence supporting or conflicting with the BN, as well as evidence justifying the exclusion of variables and relations from the BN. Although the evidence is stored within a complicated ontology structure, users can browse the

evidence in a web page that is automatically generated from the ontology, without dealing with any of the underlying technical details.

All clinicians who participated to our qualitative evaluation session indicated that clarity of evidence is an essential factor for the clinical use of prognostic models. Over 85% of the clinicians found the evidence framework a useful contribution to the current practice in prognostic models, and over 90% stated that it could be used in practice. As a next step, we plan to make a more comprehensive evaluation of the evidence framework that includes tutorials and interviews with a larger number of participants.

As further research, the evidence framework could be expanded to receive comments and suggestions from clinicians browsing the BN. A tool for auditing evidence using SPARQL could be implemented in the website generated from the ontology. This would allow clinicians from different institutions to criticise the model and suggest modifications based on clinical evidence. The evidence framework could be extended to show the strengths of evidence according to the existing evidence hierarchies (Harbour and Miller, 2001).

The evidence ontology could be proposed as a standardised format for recording BN models. Another option could be to extend the currently available XML formats for representing predictive models, such as XMLBIF (Cozman, 1998), ProbModelXML (Arias et al., 2012) and PMML (Guazzelli et al., 2009), to include information about evidence using the structure presented in this paper. The evidence-framework could also be extended to be compatible with graphical models that has bidirectional or unidirectional edges such as maximal ancestral graphs (Richardson and Spirtes 2002) or structural equation models (Pearl 2000). Such extension would offer a more flexible modelling tool for expert-driven models as the difference between causal and associational relations, and the confounding variables could be explicitly modelled.

Although clarity of evidence-base offer a potential to improve practical use of prognostic models, several other technological barriers need to be addressed to make prognostic models a part of daily clinical practice. In a recent systematic review about the impact of clinical decision support systems (CDSS), Jaspers et al. (Jaspers et al., 2011) shows that the time and effort spent for data-entry should

be minimised in order to increase the practical adoption of CDSS. In a similar review, Kawamoto et al. (Kawamoto et al., 2005) shows that CDSS should be integrated into clinical workflow and provide decision support at the time and location of decision making. Both of these factors could be achieved by integrating prognostic models into relevant clinical devices so that prognostic models could automatically extract the necessary input data from the clinical device and provide real-time predictions to clinicians.

Acknowledgments

References

Allemang, D., Hendler, J., 2011. Semantic web for the working ontologist: effective modeling in RDFS and OWL. Morgan Kaufmann.

Altman, D.G., Vergouwe, Y., Royston, P., Moons, K.G., 2009. Prognosis and prognostic research: validating a prognostic model. BMJ 338, 1432–1435.

Antal, P., Fannes, G., Timmerman, D., Moreau, Y., De Moor, B., 2004. Using literature and data to learn Bayesian networks as clinical models of ovarian tumors. Artif. Intell. Med. 30, 257–281.

Antal, P., Mészáros, T., De Moor, B., Dobrowiecki, T., 2001. Annotated Bayesian Networks: a tool to integrate textual and probabilistic medical knowledge, in: Computer-Based Medical Systems, 2001. CBMS 2001. Proceedings. 14th IEEE Symposium On. pp. 177–182.

Arias, M., Diez, F.J., Palacios-Alonso, M., Bermejo, I., 2012. ProbModelXML: A format for encoding probabilistic graphical models, in: Proceedings of the Sixth European Workshop on Probabilistic Graphical Models. Presented at the Sixth European Workshop on Probabilistic Graphical Models, Granada, Spain.

ATCBN (2015), 'Acute Traumatic Coagulopathy Bayesian Network Evidence Browser' [Online], <http://www.traumamodels.com/atcbn/ATC_EBase>, accessed 15/08/2015.

Berners-Lee, T., Hendler, J., Lassila, O., 2001. The semantic web. Sci. Am. 284, 28–37.

Bodenreider, O., 2004. The unified medical language system (UMLS): integrating biomedical terminology. Nucleic Acids Res. 32, D267–D270.

Bucci, G., Sandrucci, V., Vicario, E., 2011. Ontologies and Bayesian networks in medical diagnosis, in: System Sciences (HICSS), 2011 44th Hawaii International Conference On. pp. 1–8.

Cozman, F.G., 1998. The Interchange Format for Bayesian Networks [WWW Document]. URL http://http://www.cs.cmu.edu/~fgcozman/Research/InterchangeFormat/ (accessed 9.30.13).

Da Costa, P.C.G., Laskey, K.B., Laskey, K.J., 2008. PR-OWL: A Bayesian ontology language for the semantic web, in: Uncertainty Reasoning for the Semantic Web I. Springer, pp. 88–107.

Daly, Rónán, Shen, Qiang, and Aitken, Stuart (2011), 'Learning Bayesian networks: approaches and issues', *The knowledge engineering review,* 26 (02), 99-157.

Devitt, A., Danev, B., Matusikova, K., 2006. Ontology-driven Automatic Construction of Bayesian Networks for Telecommunication Network Management, in: 2nd Int. Workshop: Formal Ontologies Meet Industry (FOMI 2006).

Ding, Z., Peng, Y., 2004. A probabilistic extension to ontology language OWL, in: System Sciences, 2004. Proceedings of the 37th Annual Hawaii International Conference On. p. 10–pp.

Fenton, N.E., Neil, M., Caballero, J.G., 2007. Using ranked nodes to model qualitative judgments in Bayesian networks. Knowl. Data Eng. IEEE Trans. 19, 1420–1432.

Fenton, N.E., Neil, M.D., 2012. Risk Assessment and Decision Analysis with Bayesian Networks. CRC Press.

Fenz, S., 2012. An ontology-based approach for constructing Bayesian networks. Data Knowl. Eng. 73, 73–88. doi:10.1016/j.datak.2011.12.001

Guazzelli, A., Zeller, M., Lin, W.-C., Williams, G., 2009. PMML: An open standard for sharing models. R J. 1, 60–65.

Harbour, R., Miller, J., 2001. A new system for grading recommendations in evidence based guidelines. BMJ 323, 334–336.

Helsper, E.M., van der Gaag, L.C., 2002. Building Bayesian networks through ontologies, in: ECAI. p. 15th.

Helsper, E.M., van Der Gaag, L.C., 2007. Ontologies for probabilistic networks: a case study in the oesophageal-cancer domain. Knowl. Eng. Rev. 22, 67–86.

Ishak, M.B., Leray, P., Amor, N.B., 2011. Ontology-based generation of object oriented Bayesian networks, in: Proceedings of the 8th Bayesian Modelling Applications Workshop. pp. 9–17.

Jaspers, M.W., Smeulers, M., Vermeulen, H., Peute, L.W., 2011. Effects of clinical decision-support systems on practitioner performance and patient outcomes: a synthesis of high-quality systematic review findings. J. Am. Med. Inform. Assoc. 18, 327–334.

Kawamoto, K., Houlihan, C.A., Balas, E.A., Lobach, D.F., 2005. Improving clinical practice using clinical decision support systems: a systematic review of trials to identify features critical to success. Bmj 330, 765.

Knublauch, H., Fergerson, R., Noy, N., Musen, M., 2004. The Protégé OWL plugin: An open development environment for semantic web applications. Semantic Web–ISWC 2004 229–243.

Koller, D., Pfeffer, A., 1997. Object-oriented Bayesian networks, in: Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence. Morgan Kaufmann Publishers Inc., Providence, RI, pp. 302–313.

Kuo, C.-L., Buchman, D., Katiyar, A., Poole, D., 2013. Probabilistic reasoning with undefined properties in ontologically-based belief networks, in: Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence. AAAI Press, pp. 2532–2539.

Laskey, K.B., Costa, P.C., 2005. Of klingons and starships: Bayesian logic for the 23rd century, in: Uncertainty in Artificial Intelligence: Proceedings of the Twenty-first Conference.

Laskey, K.B., Mahoney, S.., 1997. Network fragments: Representing knowledge for constructing probabilistic models, in: Proceedings of the Thirteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-97). Morgan Kaufmann Publishers, Providence, RI, pp. 334–341.

Li, Junning and Wang, Z Jane (2009), 'Controlling the false discovery rate of the association/causality structure learned with the PC algorithm', *The Journal of Machine Learning Research,* 10, 475-514.

Lucas, P.J., van der Gaag, L.C., Abu-Hanna, A., 2004. Bayesian networks in biomedicine and health-care. Artif. Intell. Med. 30, 201–214.

Moons, K.G.M., Altman, D.G., Vergouwe, Y., Royston, P., 2009a. Prognosis and prognostic research: application and impact of prognostic models in clinical practice. BMJ 338, 1487–1490.

Moons, K.G.M., Royston, P., Vergouwe, Y., Grobbee, D.E., Altman, D.G., 2009b. Prognosis and prognostic research: what, why, and how? BMJ 338, b375–b375. doi:10.1136/bmj.b375

Neil, M., Fenton, N.E., Nielsen, L., 2000. Building large-scale Bayesian networks. Knowl. Eng. Rev. 15, 257–284.

Pearl, J. (2000), *Causality: models, reasoning and inference* (29: Cambridge Univ Press).

Rawlins, M., 2008. De testimonio: on the evidence for decisions about the use of therapeutic interventions. Clin. Med. 8, 579–588.

Richardson, M., Domingos, P., 2006. Markov logic networks. Mach. Learn. 62, 107–136.

Richardson, T. and Spirtes, P. (2002), 'Ancestral graph Markov models', *Annals of Statistics*, 962-1030 %@ 0090-5364.

Rieger, C., Grinberg, M., 1977. The declarative representation and procedural simulation of causality in physical mechanisms, in: Proceedings of the 5th International Joint Conference on Artificial intelligence-Volume 1. pp. 250–256.

Sadeghi, S., Barzi, A., Smith, J.W., 2005. Ontology driven construction of a knowledgebase for Bayesian decision models based on UMLS. Stud. Health Technol. Inform. 116, 223–228.

Segaran, T., Evans, C., Taylor, J., 2009. Programming the semantic web. O'Reilly Media.

Spackman, K.A., Campbell, K.E., CÃ, R., 1997. SNOMED RT: a reference terminology for health care., in: Proceedings of the AMIA Annual Fall Symposium. p. 640.

Spirtes, Peter, Glymour, Clark N, and Scheines, Richard (2000), *Causation, prediction, and search* (81: MIT press).

Tsamardinos, Ioannis and Brown, Laura E (2008), 'Bounding the False Discovery Rate in Local Bayesian Network Learning', *AAAI*, 1100-05.

Van der Gaag, L.C., Tabachneck-Schijf, H.J.M., 2010. Library-style ontologies to support varying model views. Int. J. Approx. Reason. 51, 196–208.

W3C, 2013. OWL 2 Web Ontology Language.

Wyatt, J.C., Altman, D.G., 1995. Commentary: Prognostic models: clinically useful or quickly forgotten? BMJ 311, 1539–1541. doi:10.1136/bmj.311.7019.1539

Yang, Y., Calmet, J., 2005. OntoBayes: An Ontology-Driven Uncertainty Model, in: International Conference on Computational Intelligence for Modelling, Control and Automation, 2005 and International Conference on Intelligent Agents, Web Technologies and Internet Commerce. Presented at the International Conference on Computational Intelligence for Modelling, Control and Automation, 2005 and International Conference on Intelligent Agents, Web Technologies and Internet Commerce, pp. 457–463. doi:10.1109/CIMCA.2005.1631307

Yet, B. and Marsh, D W R (2015), 'BNEvidenceBase: Evidence Framework for BN Models', <https://github.com/byet/BNEvidenceBase/. >, accessed 15/08/2015.

Yet, B., Perkins, Z., Fenton, N., Tai, N., Marsh, W., 2014a. Not just data: A method for improving prediction with knowledge. J. Biomed. Inform. 48, 28–37. doi:10.1016/j.jbi.2013.10.012

Yet, B., Perkins, Z., Tai, N., Marsh, W., 2014b. Explicit evidence for prognostic Bayesian network models. Stud. Health Technol. Inform. 205, 53–57. doi:10.3233/978-1-61499-432-9-53

Zheng, H.-T., Kang, B.-Y., Kim, H.-G., 2008. An Ontology-Based Bayesian Network Approach for Representing Uncertainty in Clinical Practice Guidelines, in: Uncertainty Reasoning for the Semantic Web I: ISWC International Workshops, URSW 2005-2007, Revised Selected and Invited Papers. Springer-Verlag, pp. 161–173.