A Knowledge Based Framework for Case-specific Diagnosis

Ganesh Ram Santhanam, Gopalakrishnan Sivaprakasam, Giora Slutzki and Samik Basu
Iowa State University, Ames, IA 50014, U.S.A.

Keywords: Diagnosis, Knowledge Bases, Description Logic, Qualitative Reasoning.

Abstract: We present a framework whereby the expert knowledge of a domain is represented as a description logic knowledge base. Based on this framework, we present an approach that uses a knowledge based system for diagnosis that allows users to key in findings for a case, and obtain the corresponding differential diagnosis for a case. The framework also prompts hypothetical findings that can effectively guide the user towards a targeted diagnosis. The framework allows iterative and interactive updates of the case specific knowledge. Computing the differential diagnosis and hypotheses can be formulated directly as conjunctive queries on the original knowledge base using the case specific knowledge. We illustrate the applicability of our framework in the context of medical diagnosis, although the approach is equally applicable in a broad range of diagnosis problems such as network forensics and criminal investigation.

1 INTRODUCTION

In many diagnostic applications, expert domain knowledge can be represented by a knowledge base (KB) that consists of a set of findings, a set of consequences, evidential relationships relating each finding to the likely consequences it indicates, and possibly other relationships between findings and diseases. Given a subset of findings that hold true in a particular 'case', *case-specific diagnosis* is the process of identifying the most likely consequences, called the *differential diagnosis* (Long et al., 1988; Miller et al., 1982). For example, in clinical diagnostic decision support (Szolovits et al., 1988; Ely et al., 2011), a doctor looks to narrow down the disease(s) that are most likely diagnoses given a set of symptoms (case) (Barnett et al., 1987; Ramnarayan et al., 2003). Similarly, the police may want to identify a set of suspects that are likely to have committed a crime given a set of supporting evidences (case) relating to the crime; and computer network administrators may want to discover suspected host(s) and the potential role(s) they could have played in launching a coordinated attack (Wang and Daniels, 2008; Pilli et al., 2010), given a set of relevant detected network events.

In this paper, we present a knowledge representation and reasoning system for *case-specific diagnosis* that guides a user progressively towards the most satisfactory (differential) diagnosis by iteratively querying the KB with respect to the evolving case. The contributions of this paper with respect to past work on diagnosis (Szolovits et al., 1988; Sohrabi et al., 2010; Baral et al., 2000) are as follows.

1. We provide a knowledge based framework for differential diagnosis where conjunctive queries are used to reason about domain knowledge that is relevant to a given case, taking into consideration case-specific evolution of findings.
2. We allow case-specific knowledge to evolve by the addition of findings (e.g., a patient may present new symptoms), and treat case-specific diagnosis as an ***iterative*** process of recomputing and refining the differential diagnosis based on new test results relevant to the case.
3. Moreover, we treat case-specific diagnosis as an ***interactive*** process, i.e., a user can change the case-specific knowledge by adding different sets of findings during the course of diagnosis, leading to the corresponding changes in the differential diagnoses in the subsequent iteration.

Organization. The rest of the paper is organized as follows. Section 2 describes a medical diagnosis tool called Dxplain used to demonstrate our framework, and how domain knowledge can be represented as a Description Logic (DL) knowledge base. Section 3 describes how we reason with domain knowledge using DL conjunctive queries. Section 4 discusses how conjunctive queries can be extended to compute suggested hypotheses during diagnosis. Section 5 summarizes and discusses directions for future research.

Ram Santhanam G., Sivaprakasam G., Slutzki G. and Basu S..
A Knowledge Based Framework for Case-specific Diagnosis.
DOI: 10.5220/0005272104820488
In *Proceedings of the International Conference on Agents and Artificial Intelligence* (ICAART-2015), pages 482-488
ISBN: 978-989-758-074-1

2 BACKGROUND

We present our framework in the context of medical diagnosis, where the basic knowledge includes facts about the findings/symptoms, the diseases and their relationships.

2.1 Case Study: Medical Diagnosis

Dxplain (Barnett et al., 1987; Hoffer et al., 2005) is an expert system for clinical diagnostic decision support developed and enhanced by the Massachusetts General Hospital since 1984. The Dxplain KB consists of 5000 clinical findings and 2400 diseases and their relationships (averaging about 50 findings per disease). In the 'case analysis mode,' DXplain uses a graphical interface to collect findings of a case from a user (doctor) and produces a list of diseases that are likely to be the consequences of the given findings. DXplain also provides for each disease a set of findings that support that disease at various levels of intensity. Dxplain uses Bayesian probabilistic estimates of likelihood in order to quantify the strength of evidence of a finding to a disease and rank diseases. In contrast, the proposed framework will allow users to address different types of questions using the DL query constructs, thus broadening the scope of differential diagnosis beyond what is achievable using Dxplain.

2.2 Semantic Information in Dxplain

Findings. A finding is a clinical manifestation (clinical symptoms, signs, epidemiological data, personal biodata) or the result of laboratory tests. Each finding has an attribute *'Finding Importance'* that can hold values on a totally ordered (ordinal) scale. This attribute indicates how important it is to explain the presence of the finding (considering its significance). For example, a *runny nose* is less important than *acute right abdominal pain*.

Diseases. Each disease is described by three attributes.

1. *Disease Importance*, indicates the rating on a totally ordered value scale from 1 to 5 (higher values signifying higher importance).
2. *Disease Prevalence*, indicates how common the disease is on a totally ordered (ordinal) scale: Very Common, Common, Rare and Very Rare.
3. *Disease Urgency*, indicates whether urgent action is needed or not, if the disease was actually present.

Table 1: Attributes and their Ordinal Scales.

Attribute	Label	Range
Finding Importance	η_f	$\{\eta_1^f, \eta_2^f, \eta_3^f, \eta_4^f, \eta_5^f\}$
Disease Importance	η_d	$\{\eta_1^d, \eta_2^d, \eta_3^d, \eta_4^d, \eta_5^d\}$
Disease Prevalence	ϕ_d	$\{\phi_1^d, \phi_2^d, \phi_3^d, \phi_4^d\}$
Disease Urgency	υ_d	$\{Urgent, Normal\}$
Finding Frequency	μ_{fd}	$\{Indicates, Eliminates\}$
Strength of Support	σ_{fd}	$\{\sigma_{-1}, \sigma_0, \sigma_1, \sigma_2\}$

Finding - Disease Relationships. A relationship between a finding and a disease signifies that the finding either increases or diminishes the support (evidence) for the existence of that disease. This relationship is described by two attributes.

1. *Finding Frequency/Correlation*, indicates how frequently the finding is seen in the disease, or whether it rules out the disease. This attribute is evaluated on a qualitative scale '*indicates*' to denote that the disease should always be considered given the finding; and '*eliminates*' to denote that the presence of the finding excludes the disease.
2. *Strength of Support*, indicates how strongly the finding suggests or supports the disease. This is evaluated on a qualitative (ordinal) scale (Supports / Strongly supports / Very strongly supports; or Supports Absence[1]).

Disease - Disease Relationship: Similarity. For each disease d, Dxplain maintains a set of similar diseases that should be considered if d is being considered in the diagnosis.

The attributes of the findings, diseases and their relationships, and their values on an ordinal (rather than cardinal) scale are defined in Table 1.

2.3 Knowledge Base Example

We illustrate and motivate representation and reasoning about domain and case-specific knowledge in the proposed system using a simplified example of a medical KB[2] with 5 findings and 4 diseases. The graph representation of the KB as well as the evolution of the case in terms of the findings of a patient entered by a doctor is shown in Figure 1. Note that f_1 supports, and f_3 supports the absence of d_1; f_2 strongly supports, f_3 very strongly supports, and f_4 supports the absence of d_2; f_4 supports d_3; f_5 very strongly

[1] Supports the absence of a disease.

[2] The KB in this example is simplified compared to the actual medical KB in Dxplain or other tools such as Isabel.

supports d_4; and d_3 and d_4 are similar to each other. The sequence of steps taken by the doctor (D) and the system (S) in tandem during diagnosis is as follows.

- [D:] *Provide Initial Case Findings.* Suppose the doctor provides the initial case findings for a patient, say $\{f_1, f_2\}$.
- [S:] The system then performs two computations:
  - *Differential Diagnosis.* The system computes the differential diagnosis $\{d_1, d_2\}$ because f_1 and f_2 support d_1 and d_2 respectively. Note that d_1 and d_2 cannot be distinguished based on the available evidence, as the level of support offered by the respective findings in the case is σ_1.
  - *Hypothesis Generation.* Further, the system suggests a set $\{f_3, f_4\}$ of hypothetical findings back to the doctor. The finding f_3 is suggested because it enhances the support for d_2 while diminishing support for d_1 and f_4 is suggested because it diminishes the support for d_2, leaving the support for d_1 unchanged.
- [D:] *Update of Findings.* On being presented with the differential diagnosis and the corresponding hypotheses, the doctor may choose to test (by a physical exam, ordering a laboratory procedure such as a blood test, etc.) the presence or absence of some of them based on his own expertise or experience. When the results are known, the doctor may update[3] the case findings by adding one or more of the suggested hypothesis that tested positive for the patient and remove the ones for which the patient tested negative.
- [S:] *Recomputation of Differential Diagnosis.* Consider the scenario, where the doctor adds f_3 (as its presence is validated). This results in the case scenario presented in the bottom-left of Figure 1. Note that the addition of f_3 adds no new diseases, but changes the level of support to d_1 and d_2 in a way that allows the doctor to infer that d_2 is more likely than d_1. On the other hand, if f_4 were added to the case, then it can lead to a differential diagnosis by which d_1 may be inferred to be more likely than d_2, in addition to the inclusion of d_3 in the differential diagnosis. Finally, if both f_3 and f_4 are added to the case, then the resulting differential diagnosis includes d_3, and the the doctor may infer that d_2 is more likely than d_1 given the evidence.

[3]While the hypotheses only provide suggestions for tests that the doctor may order, updates are based on results of such tests.

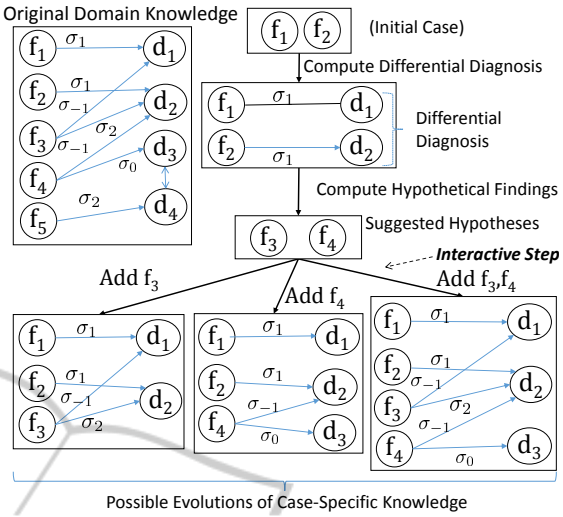

Figure 1: Original KB and Case Analysis Example.

In the above example, at each step the addition of new findings results in computation of (a) the new differential diagnosis; and (b) a set of hypothetical findings relevant to the case. Note that the doctor can choose various subsets of the set of hypotheses presented to him in the case, each leading to possibly different differential diagnoses in the successive steps. In the following section, we present the DL KB for capturing the domain-specific and case-specific knowledge followed by the queries to the knowledge base that computes the differential diagnosis and hypotheses.

3 REASONING WITH KB USING DLs

We assume basic knowledge of DLs (Baader et al., 2003) for the rest of the paper. A DL KB is a tuple $\mathcal{K} = \langle \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$ where $\mathcal{A}$ is an ABox, $\mathcal{T}$ is a TBox and $\mathcal{R}$ is an RBox. The TBox and RBox in $\mathcal{K}$ constitute the schema or structure underlying the domain knowledge and the ABox specifically contains knowledge about individuals in $\mathcal{K}$. $\mathcal{A}$ is a set of concept and role assertions of the form $C(a)$ and $r(a, b)$ (unary and binary predicates) respectively, where C is a concept, r is a role, and a and b are individuals. $\mathcal{T}$ and $\mathcal{R}$ are sets of terminological axioms that state how concepts and roles are related to each other (respectively).

3.1 Diagnostic Knowledge as a DL KB

We now illustrate how to model the domain knowledge in terms of a DL KB $\mathcal{K} = \langle \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$ for our

running example from the domain of medical diagnosis. The concept and role assertions in the ABox $\mathcal{A}$ and the list of concept and role inclusions in $\mathcal{T}$ and $\mathcal{R}$ are given in Figure 2.

Concepts and Roles for Domain Knowledge. The concept names φ and δ in $\mathcal{K}$ correspond to individual findings and diseases respectively. The role names σ_{-1}, σ_0, σ_1 and σ_2 correspond to relationships between findings and diseases in the original domain knowledge with the respective evidential strengths indicated by the corresponding subscripts. These constitute generic domain knowledge common to all the conceivable cases to be analyzed.

Concepts to Encode Attributes of Findings and Diseases. The attributes of findings and diseases, including finding and disease importance, finding frequency, disease urgency, etc. are encoded as general concepts in $\mathcal{K}$. For example, $\mathcal{K}$ will have concepts $\eta_1^d, \ldots, \eta_5^d$, $\eta_1^f, \ldots, \eta_5^f$, μ_f, υ_d, etc., corresponding to the labeling functions (see Table 1).

Concept Hierarchy. Attributes such as finding and disease importance have different levels of valuations in Dxplain that are totally ordered. We model one general concept for each valuation of that attribute and whenever appropriate, impose the total ordering over these valuations using concept inclusions. For example, disease importance is an attribute that has valuations from 1 to 5, and we model this with inclusions $\eta_i^d \sqsubseteq \eta_{i+1}^d$ for $i = 1, \ldots, 4$ in $\mathcal{T}$.

Role Hierarchy. Roles such as σ_1, σ_2, etc. indicate various degrees of relationship, e.g., strength of support of a disease by a finding. We model the fact that one role represents a higher degree of support than another by role hierarchies. For example, to model the fact that a finding f which very strongly supports a disease d also strongly supports it, we include a role inclusion axiom $\sigma_2 \sqsubseteq \sigma_1$ in the RBox $\mathcal{R}$. Note that σ_{-1} is a role that denotes that the finding offers evidence in support of the absence of the disease.

We assume that assertions in $\mathcal{K}$ pertaining to the above concepts and roles are not updated during the course of case analysis in Dxplain, as they constitute domain knowledge that is not expected to change during the course of the diagnosis.

Concepts for Case-specific Knowledge. As mentioned above, a *case* is a set of findings, that may include (a) Positive Findings that are marked as present in the case; (b) Negative Findings that are marked to be absent in the case; and (c) Unknown Findings that are marked as being of no interest to the case. The concept names φ_+, φ_- and φ_u correspond to

$\mathcal{A}$

$\varphi^{\mathcal{I}} = \{f_1, f_2, f_3, f_4, f_5\}$
$\delta^{\mathcal{I}} = \{d_1, d_2, d_3, d_4\}$
$\sigma_{-1}^{\mathcal{I}} = \{(f_3, d_1), (f_4, d_2)\}$
$\sigma_0^{\mathcal{I}} = \{(f_4, d_3)\}$
$\sigma_1^{\mathcal{I}} = \{(f_1, d_1), (f_2, d_2)\}$
$\sigma_2^{\mathcal{I}} = \{(f_3, d_2), (f_5, d_4)\}$
$\eta_1^{f\mathcal{I}} = \{f_1, f_2\} \ldots$
$\eta_1^{d\mathcal{I}} = \{d_2, d_4\} \ldots$
$\phi_1^{d\mathcal{I}} = \{d_1, d_2\} \ldots$
$\upsilon^{d\mathcal{I}} = \{d_4, d_2\} \ldots$

$\mathcal{T}$

$\eta_1^f \sqsubseteq \eta_2^f \cdots \sqsubseteq \eta_5^f$
$\eta_1^d \sqsubseteq \eta_2^d \cdots \sqsubseteq \eta_5^d$
$\phi_1^d \sqsubseteq \phi_2^d \cdots \sqsubseteq \phi_4^d$
$\varphi_c = \varphi_+ \sqcup \varphi_- \sqcup \varphi_u$
$\varphi_c \sqsubseteq \varphi$
$\delta_c \sqsubseteq \delta$

$\mathcal{R}$

$\sigma_2 \sqsubseteq \sigma_1 \sqsubseteq \sigma_0 \sqsubseteq \sigma$
$\sigma_{-1} \sqsubseteq \neg\sigma_0 \sqsubseteq \sigma$

Figure 2: DL Encoding of Original Domain KB $\mathcal{K}$.

the positive, negative and unknown case-specific findings. For example, during case analysis, when a doctor inputs a finding as present (absent or unknown) for the current case, the corresponding finding in $\mathcal{K}$ is labeled with the concept φ_+ (φ_- or φ_u respectively). We also define a concept $\varphi_c = \varphi_+ \sqcup \varphi_- \sqcup \varphi_u$ that corresponds to the set of all findings (positive, negative or unknown) in the case.

3.2 Case-Specific Diagnosis

A new case is initiated by the doctor by keying in a set C of positive and negative findings. The system then queries the KB to obtain (a) the differential diagnosis $\mathcal{D}(C)$and (b) a set of *suggested findings* or *hypotheses* denoted by $\mathcal{H}(C)$ that are not already in the case (and possibly could help refine the diagnosis by increasing or decreasing the evidential support of, or eliminating one or more diseases in the differential).

Conjunctive Queries for DL Reasoning. In order to compute $\mathcal{D}(C)$ and $\mathcal{H}(C)$ in each iteration, we construct appropriate conjunctive queries on $\mathcal{K}$. A *conjunctive query* (CQ) (Glimm et al., 2008) is an existentially quantified conjunction of atoms on a set of variables V such that each atom asserts a concept membership or role relationship between one or more of the variables (or individuals). A CQ can optionally have a free variable (not in V), certain valuations of which may satisfy the CQ. For example, $\exists y, z.Q(x) \wedge P(x, y) \wedge P(y, z) \wedge P(z, x)$ is a CQ where $V = \{y, z\}$, x is a free variable, and Q and P correspond to a concept and a role respectively. The CQ

$$P(d) = \exists f.\ \varphi_c(f) \wedge \delta(d) \wedge \sigma_0(f, d)$$

states that there exists a finding in the case C that supports some disease d in the domain. The variable d is free in the query, resulting in a set $\mathcal{D}(C)$ representing

the differential diagnosis. Note that, in view of the role hierarchy specified in $\mathcal{R}$, the last conjunct σ_0 accounts for support offered by f to d at any level. Once $P(d)$ is computed, the system adds each such disease d to the case-specific knowledge using the concept δ_c, i.e., an assertion $\delta_c(d)$ is added to the case-specific ABox. We next formulate the generation of hypotheses as a CQ that obtains all findings in $\mathcal{K}$ that provide either positive support of support of the absence of any of the existing diseases in the differential diagnosis $\mathcal{D}(C)$.

$$Q(f) = \exists d.\, \neg\varphi_c(f) \wedge \delta_c(d) \wedge \sigma(f, d)$$

The above CQ states that there exists a disease d in the differential diagnosis $\mathcal{D}(C)$ that is offered some evidential support by a finding f that is not present in the case findings C. Note that the role σ is defined in $\mathcal{R}$ as one that includes all types of evidential support (including support of the absence of disease) from findings and diseases. Since f is a free variable, $Q(f)$ obtains a set of findings exactly corresponding to those in $\mathcal{H}(C)$. Also, the negation in the first term (checking that f is not in the case-specific knwoledge so far) can be viewed as a syntactic construct, and hence will not make query answering undecidable.

Once a disease differential $\mathcal{D}(C)$ and the set $\mathcal{H}(C)$ of suggested findings are presented to the user, the user may chose to proceed by changing the case in one of two ways: (a) by marking one or more of the suggested findings as Positive, Negative or Unknown, or (b) by including findings that are not in $\mathcal{H}(C)$ as Positive or Negative for the case. This triggers a new iteration for computing $\mathcal{D}(C')$ and $\mathcal{H}(C')$ with respect to new case specific findings C'. The user continues to change the case iteratively until he/she is satisfied with the obtained disease differential (i.e., the diagnosis is satisfactory).

We represent the initial set of findings for a case by the ABox $\mathcal{C}_0$. Further, we view iterative addition of findings to the case as a sequence of ABoxes beginning with $\mathcal{C}_0$, namely $\mathcal{C}_1, \mathcal{C}_2, \ldots, \mathcal{C}_n$. $\mathcal{C}_n$ is an ABox that provides a differential diagnosis that satisfies a target criteria defined by the user. For instance, the user may want to include at least one serious disease with very high level of evidence from case findings.

Figure 3 illustrates a possible evolution of ABoxes for the running example in Figure 1. The evolution corresponds to the successive addition of a set of positive findings to the initial case. Each ABox contains assertions (δ_c) corresponding to the diseases that are included in the differential diagnosis for the current set of case findings (φ_+). For readability, each ABox in Figure 3 shows only assertions that correspond to case-specific knowledge during the iteration.

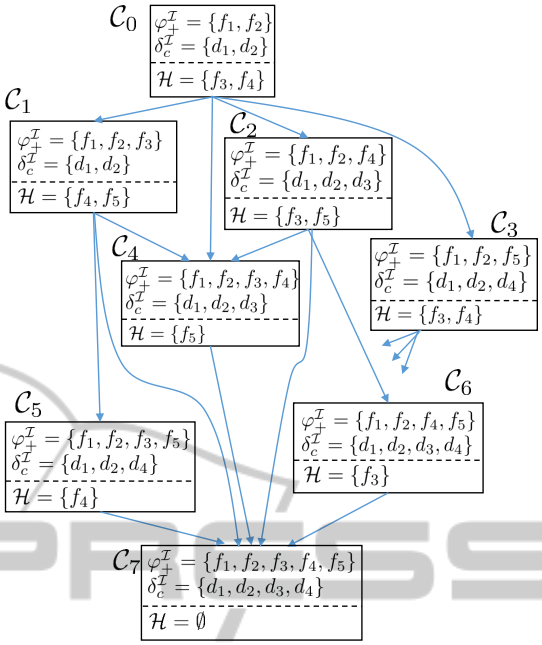

Figure 3: Evolution of ABoxes during Case Analysis for the example in Figure 1. The contents of the ABoxes are shown in terms of the interpretation of concepts rather than the assertions (e.g., $\delta_c(d_1), \delta_c(d_2)$) themselves. Each ABox is also annotated with a set $\mathcal{H}$ of hypotheses.

4 FINDING TARGETED HYPOTHESIS

The goal of providing suggested hypotheses to the doctor is to prompt the doctor about case information (findings), whose inclusion into the case may alter the differential diagnosis in desired ways. In the previous section, we presented one way to compute $\mathcal{H}(C)$ as the set of findings that either increase or decrease the evidential support for one or more diseases in $\mathcal{D}(C)$. There may be other potentially useful hypotheses that help the doctor navigate the space of differential diagnoses as desired, during case evolution; here we present three such examples.

a. A doctor may have a strong suspicion that one of the diseases, $d \in \mathcal{D}(C)$ is actually present in the patient. In such a case, the doctor may find it useful if prompted (only) with the findings that are known to increase the evidence for d (based on the original domain knowledge). Given a disease d, the following CQ obtains findings f (free variable) that are not already in the case C, which also very strongly[4] support disease d.

[4]Note that this is only one possible way to obtain a set of findings that strongly support a specific disease d.

$$Q'_d(f) = \neg\varphi_c(f) \ \wedge \ \delta_c(d) \ \wedge \ \sigma_2(f,d)$$

b. A doctor may choose to first expand the differential diagnosis to include diseases that are of high importance level (e.g., require immediate attention) and are possible given the case. Recall that such diseases are labeled with the concept η_5^d in $\mathcal{K}$. We compute the desired hypotheses in this case by looking for findings that support both a $d \in \mathcal{D}(C)$ as well as a $d' \notin \mathcal{D}(C)$ such that $\eta^d(d') = \eta_5^d$. The corresponding CQ is as follows.

$$Q'(f) = \exists d, d'.\big[\ \neg\varphi_c(f) \ \wedge \ \delta_c(d) \ \wedge \ \neg\delta_c(d') \ \wedge \ \eta_5^d(d') \wedge \ \sigma_0(f,d')\big]$$

c. A doctor may want to include diseases of importance level at least 4 and are similar to any of the diseases presently in the differential diagnosis. The corresponding set of findings is then obtained by the following CQ, assuming ψ represents the *similarity* relation between diseases.

$$Q'(f) = \exists d, d'.\big[\ \neg\varphi_c(f) \ \wedge \ \delta_c(d) \ \wedge \ \neg\delta_c(d') \ \wedge \ \eta_4^d(d') \wedge \ \psi(d,d') \ \wedge \ \sigma_0(f,d')\big]$$

Note that in the above CQ, the diseases of importance level 4 and 5 will be included because of the inclusion axiom $\eta_5^d \sqsubseteq \eta_4^d$ in the TBox.

5 SUMMARY & FUTURE WORK

In this paper, we present a knowledge based framework for addressing the case-specific diagnosis problem. The framework allows users to obtain the differential diagnosis for a case, and prompts hypothetical findings that can effectively guide the user towards a diagnosis that is supported by evidences from case findings. The framework facilitates iterative and interactive updates of case specific knowledge as an evolution of sequences of ABoxes. We show that the queries relevant to differential diagnosis and hypotheses generation can be formulated directly as conjunctive queries on the original knowledge base using the case specific knowledge. We present the applicability of our framework in the context of medical diagnosis. We note that however, the **generality** of our framework makes it applicable to other diagnosis problems such as network forensics and criminal investigation.

In addition to its generality, the proposed framework provides an approach for addressing relevant and interesting problems in diagnosis. One of the important requirements of any diagnostic system is the **justification** of diagnosis, i.e., what portions of the domain knowledge and which findings can be used to explain a conclusion for a particular instance of case-specific diagnosis. We can facilitate justification by computing a proof of correctness of the results obtained from the conjunctive queries executed.

Finally, we conjecture that the interactive and iterative nature of our framework allows for effective diagnosis and discovery of consequences that were previously unknown; this is likely to have significant impact in medical domain resulting in **discovery** of new relationships between findings and diseases, as also in identifying new traits of diseases. In this sense, this paper provides a road map for addressing various challenging problems in diagnosis.

ACKNOWLEDGMENTS

We are grateful to Edward P. Hoffer, M.D., Associate Clinical Professor of Medicine at the Harvard Medical School for clarifying to us the working of Dxplain and Jason Maude, CEO and Co-founder, Isabel Healthcare for his insightful comments on the diagnostic tool Isabel. The first three authors were partially supported by NSF grant 1116050.

REFERENCES

Baader, F., Calvanese, D., McGuinness, D. L., Nardi, D., and Patel-Schneider, P. F., editors (2003). *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press.

Baral, C., McIlraith, S., and Son, T. (2000). Formulating diagnostic problem solving using an action language with narratives and sensing. In *Proceedings of the 7th International Conference on Principles of Knowledge Representation and Reasoning*, pages 311–322.

Barnett, G. O., Cimino, J. J., Hupp, J. A., and Hoffer, E. P. (1987). Dxplain: an evolving diagnostic decision-support system. *Journal of the American Medical Association*, 258(1):67–74.

Ely, J. W., Graber, M. L., and Croskerry, P. (2011). Checklists to reduce diagnostic errors. *Academic Medicine*, 86:1040–2446.

Glimm, B., Horrocks, I., Lutz, C., and Sattler, U. (2008). Conjunctive query answering for the description logic. *Journal of artificial intelligence research*, 31:157–204.

Hoffer, E. P., Feldman, M. J., Kim, R. J., Famiglietti, K. T., and Barnett, G. O. (2005). Dxplain: patterns of use of a mature expert system. In *Proceedings of the Annual Symposium of the American Medical Informatics Association*, volume 2005, pages 321–325.

Long, W. J., Naimi, S., Criscitiello, M., and Larsen, G. (1988). Differential diagnosis generation from a causal network with probabilities. In *Proceedings of*

the International Conference on Computers in Cardiology, pages 185–188. IEEE.

Miller, R. A., Pople Jr, H. E., and Myers, J. D. (1982). Internist-I, an experimental computer-based diagnostic consultant for general internal medicine. *New England Journal of Medicine*, 307:468–476.

Pilli, E. S., Joshi, R. C., and Niyogi, R. (2010). Network forensic frameworks: Survey and research challenges. *Digital Investigation*, 7(1):14–27.

Ramnarayan, P., Tomlinson, A., Rao, A., Coren, M., Winrow, A., and Britto, J. (2003). Isabel: a web-based differential diagnostic aid for paediatrics: results from an initial performance evaluation. *Archives of disease in childhood*, 88(5):408–413.

Sohrabi, S., Baier, J. A., and McIlraith, S. A. (2010). Diagnosis as planning revisited. In *Proceedings of the 12th International Conference on the Principles of Knowledge Representation and Reasoning*, pages 26–36.

Szolovits, P., Patil, R. s., and Schwartz, W. b. (1988). Artificial intelligence in medical diagnosis. *Annals of Internal Medicine*, 108(1):80–87.

Wang, W. and Daniels, T. E. (2008). A graph based approach toward network forensics analysis. *ACM Transactions on Information and System Security*, 12(1):4.