

Abductive Reasoning with Description Logics: Use Case in Medical Diagnosis

Júlia Pukancová and Martin Homola

Comenius University in Bratislava,
Mlynská dolina, 84248 Bratislava, Slovakia
{pukancova, homo1a}@fmph.uniba.sk

Abstract. Ontologies have been increasingly used as a core representation formalism in medical information systems. Diagnosis is one of the highly relevant reasoning problems in this domain. In recent years this problem has captured attention also in the description logics community and various proposals on formalising abductive reasoning problems and their computational support appeared. In this paper, we focus on a practical diagnostic problem from a medical domain – the diagnosis of diabetes mellitus – and we try to formalize it in DL in such a way that the expected diagnoses are abductively derived. Our aim in this work is to analyze abductive reasoning in DL from a practical perspective, considering more complex cases than trivial examples typically considered by the theory- or algorithm-centered literature, and to evaluate the expressivity as well as the particular formulation of the abductive reasoning problem needed to capture medical diagnosis.

Keywords: Diagnosis, abduction, use case.

1 Introduction

Abduction, originally introduced by Peirce [15], is a form of backward reasoning, typically with a diagnostic rationale. We have a knowledge base \mathcal{K} that is supposed to model some problem, and we have an observation O which is supposed to follow in situations captured by \mathcal{K} , but we are not able to explain O deductively, i.e., $\mathcal{K} \not\models O$. In abductive reasoning we ask the question – why is it that O does not follow from \mathcal{K} , and we look for a hypothesis (or, explanation) H such that, if added to \mathcal{K} , then O will follow from the resulting knowledge base. Moreover, we most typically look for explanations consisting of extensional rather than intensional knowledge, i.e., some set of ground facts that will, together with \mathcal{K} , explain O .

Abduction only recently captured the researchers' interest also in the area of ontologies and DL [7], where it also has some interesting applications, including possible explanations of incomplete modelling or incomplete matching [4], monitoring malfunctions in complex systems [11], and multimedia interpretation [16], among others.

The problem of diagnosis, often shown as a classic example of abductive reasoning [7,9], is highly relevant in the medical domain. Not only for primary diagnosis of a certain disease (as the model example in this use case), but also in emerging applications such as telemedical monitoring systems and ambient assisted living, where the patient's

condition is continually monitored and diagnosed for anomalies, therapy adherence, etc. Most of these applications nowadays heavily rely on ontologies (i.e., DL-based knowledge bases) that have been increasingly used as core representation formalism for clinical knowledge. While abduction over DL has been studied especially from the theoretical and from the algorithmic perspective, we are not aware of any case studies focusing on the practical aspects of modelling problems for abductive reasoning.

In this paper, we focus on a practical diagnostic problem from a medical domain: the diagnosis of diabetes mellitus. Based on information from clinical guidelines and other relevant sources (e.g., [1,2]) we formalize it in DL in such a way that the expected diagnoses are abductively derived. While we simplify the problem for reasons of conciseness, we do abstract a number of distinct, less or more problematic cases that need to be addressed, including: (a) dealing with the hierarchy of symptoms and possible diagnoses, (b) differential and elimination diagnosis, (c) associated conditions with similar symptoms, (d) distinguishing and reporting complications, and some more.

In the analysis that follows, we evaluate the modelled examples from the perspective of which particular variant of abduction is being addressed, what DL expressivity is needed, and we highlight the most important modelling issues that we run into.

In the end, we learned the following lessons: medical diagnosis especially requires ABox abduction, as hypothesizing the intensional knowledge in this domain is typically not desired – that is the area of domain experts. We were mostly able to model our simplified examples with the rather less expressive DL \mathcal{ALC} for which abductive reasoning is available [9,12,13]. Though, examples requiring more complex constructs can also be found. Finally, modelling diagnostic knowledge bases with DL is different from modelling typical ontologies. In order to get the desired explanations the statements often need to be formulated more strongly, so that the desired observations follow. Also, to compare the generated hypotheses, at least part of the knowledge is used also deductively. Combining abductive and deductive reasoning within one knowledge base poses some difficulties, even on simplistic examples.

2 Abductive Reasoning with DL

The basic abductive framework for DL was introduced by Elsenbroich et al. [7] who proposed formulations for a number of distinct abductive problems. The main three types are summarized below. We will assume that the reader is already familiar with DL, the split of the knowledge base into the TBox (intensional knowledge) and the ABox (extensional knowledge), basic syntax, and semantics [3].

Definition 1 (Abduction problems in DL [7]). *An abduction problem is a pair $\mathcal{P} = (\mathcal{K}, O)$ such that \mathcal{K} is a knowledge base in DL, and O a TBox or ABox assertion. A solution of \mathcal{P} is any finite set H of TBox and ABox assertions such that $\mathcal{K} \cup H$ is consistent and $\mathcal{K} \cup H \models O$. In addition, \mathcal{P} is called*

- **TBox abduction problem:** if H is a set of TBox assertions and O is a TBox assertion.
- **ABox abduction problem:** if H is a set of ABox assertions and O is an ABox assertion.

- **Knowledge base abduction problem:** the general problem, i.e., if there are no restrictions on H and O .

The definition gives a generic framework for abduction, but it is not very useful without further constraining the possible hypotheses. The number of possible explanations is very high, even infinite, therefore we need to be able to compare them and select the most preferred ones. The commonly used restrictions include [7]:

Definition 2. Given an abduction problem $\mathcal{P} = (\mathcal{K}, O)$ and hypotheses H, H' , we say that:

1. H is **consistent** if $H \cup \mathcal{K} \not\models \perp$, i.e. H is consistent w.r.t. \mathcal{K}
2. H is **relevant** if $H \not\models O$, i.e. H does not entail O
3. H is **explanatory** if $\mathcal{K} \not\models O$, i.e. \mathcal{K} does not entail O
4. H is **stronger** than H' ($H <_{\mathcal{K}} H'$) if $\mathcal{K} \cup H$ entails H' (and vice-versa H is **weaker** than H' if $\mathcal{K} \cup H'$ entails H)
5. **minimal** if for every H' , $H' <_{\mathcal{K}} H$, i.e., H is weaker than any other H'

In general, H is a preferred solution if it is consistent, relevant, explanatory and there is no strictly weaker solution H' (i.e., such that $H <_{\mathcal{K}} H'$ and $H' \not<_{\mathcal{K}} H$). If there is single such solution, it is called the most preferred. Such hypotheses are (semantically) minimal. Minimality is important, because if there is a (strictly) weaker hypothesis than H it means that H hypothesizes too much. As abduction amounts to guessing, in a sense we do not want to guess more than necessary in order to derive the observation.

Consistency is required, because from inconsistency ($\mathcal{K} \cup H \models \perp$) one is able to derive everything. Such hypotheses would explain every observation and so it is not meaningful to find solutions not consistent with the knowledge base \mathcal{K} . The knowledge base \mathcal{K} represents the background theory from which we are interested to derive the hypotheses. Therefore we should not be able to explain the observation without it. Such hypotheses (i.e., when $H \models O$) are therefore irrelevant. Similarly, an abduction problem only needs explaining the observation does not already follow from \mathcal{K} .

The diagnostic problems in the medical domain, which is our interest in this paper, will most often call for ABox abduction. This is because our aim is not to enrich the knowledge base with new axioms; these are typically sufficiently described by domain experts. Therefore our purpose is to build a knowledge base formalizing the available expert's knowledge (TBox) and use it to find explanations in form of facts (ABox assertions) to any observation, which is typically also a fact (ABox assertion).

A number of researchers addressed the computational solution for abduction problems. We will focus on those who addressed ABox abduction. Klarman et al. [12] proposed an algorithm for ABox abduction on top of \mathcal{ALC} based on resolution. They show that it is sound and complete. Halland and Britz [9] developed, also for \mathcal{ALC} , a method based on the DL tableau algorithm. Ma et al. [13] also rely on the DL tableau algorithm, but extend the approach towards \mathcal{ALCI} . Completeness was not shown by Ma et al., while Halland and Britz explicitly note that their approach is incomplete.

Du et al. [5] solve abduction reasoning in DL via a reduction to logic programming, where abduction has been extensively studied [6]. In logic programming, abductive explanations are not arbitrary, but are typically drawn from a set of distinctive literals

called abducibles. This is because the user is often able to characterize in which part of the knowledge the hypothesis is expected, and thus to reduce the search space. To transfer this notion to the area of DL, Du et al. introduced a new variant of the ABox abduction problem, which we will call *simple*, in the form of $\mathcal{P}(\mathcal{K}, A, O)$. Besides for the knowledge base \mathcal{K} and the observation O (set of concept assertions or atomic roles assertions), it adds the abducibles A – a set of atomic concepts and atomic roles. An abductive solution H for $\mathcal{P} = (\mathcal{K}, A, O)$ is a minimal set of ABox axioms composed of individuals of \mathcal{K} and concepts or roles of A , such that $\mathcal{K} \cup H \models O$. The solution should be relevant, and consistent. The simple ABox abduction problem may be solved by reduction to logic programming, and consecutive evaluation by a readymade reasoning engine for logic programming.

3 Diabetes Mellitus Use Case

of the assumed simplification

Diabetes mellitus (DM) is a group of metabolic diseases characterized by hyperglycemia resulting from defects in insulin secretion, insulin action, or both. The chronic hyperglycemia of diabetes is associated with long-term damage, dysfunction, and failure of various organs, especially the eyes, kidneys, nerves, heart, and blood vessels [2]. We chose DM for our use case, as its diagnosis is a complex problem, with need to distinguish between particular subtypes and associated conditions, identification of possible complications, etc.

Our aim is to conceptualize a KB in DL that can be used for diagnosis of DM relying on abductive reasoning. As the problem is rather complex, we will concentrate on selected specific subproblems, and we will also abstract from some details which can be implemented analogously.

3.1 Hierarchy of Symptoms

Typical symptoms of diabetes mellitus include: frequent urination, excessive thirst, hyperglycemia, blurred vision, anorexia, weight loss, fatigue, and weakness. There are some more, but they can be added into the formalization analogously. In the following, diabetes mellitus is represented by the concept symbol DM, and the symptoms by S_1, \dots, S_8 , respectively. In abductive reasoning, we observe some set of symptoms and try to hypothesize the most relevant diagnosis. Therefore we record the relation between the diagnosis and its symptoms by the axiom:

$$\exists \text{hasDiag.DM} \sqsubseteq \exists \text{hasSymp.}(S_1 \sqcap \dots \sqcap S_8) \quad (1)$$

While literally (more precisely: deductively) this axiom means, that whoever has DM must simultaneously manifest all eight symptoms S_1, \dots, S_8 , which is not always true; it allows to generate relevant abductive hypotheses: if the patient p is observed to have any subset of the symptoms S_1, \dots, S_8 (e.g., we have an observation $O_1 = p: (\exists \text{hasSymp.S}_2) \sqcap (\exists \text{hasSymp.S}_8)$) then $H_1 = \{p: \exists \text{hasDiag.DM}\}$ is among the generated diagnoses. Note that this kind of modelling is simplified, as it ignores uncertainties, often present in medical knowledge. In this paper we take this simplification as

we want to explore the possibilities of abduction with regular DL. Introducing uncertainties is left for future work.

There are, possibly, some relations between symptoms, like one symptom may be a more specific or a synonymous name for another, for instance, polydipsia (S_9) is a synonym for excessive thirst. This is modelled by adding:

$$S_2 \equiv S_9 \quad (2)$$

Now, whenever we observe S_9 instead of S_2 among some other symptoms of DM we derive equal hypotheses (e.g., if $O_2 = \{p: (\exists \text{hasSymp}.S_9) \sqcap (\exists \text{hasSymp}.S_7)\}$, H_1 is still abductively derived).

A more complex relationship among symptoms may occur in cases when some symptoms are conditions which may themselves be manifested by some other symptoms. For instance anorexia (S_5) is a condition associated with weight loss, fatigue, and weakness (S_6, \dots, S_8). This may be modelled in two ways, adding either (3–4) or (5):

$$\begin{aligned} S_5 \sqsubseteq \exists \text{hasSymp}.(S_6 \sqcap \dots \sqcap S_8) & \quad (3) \\ \text{hasSymp} \cdot \text{hasSymp} \sqsubseteq \text{hasSymp} & \quad (4) \end{aligned} \qquad S_5 \sqsubseteq S_6 \sqcap \dots \sqcap S_8 \quad (5)$$

This then allows us to keep the axioms that relate diagnoses to symptoms more concise, e.g., we may now replace (1) by:

$$\exists \text{hasDiag}.DM \sqsubseteq \exists \text{hasSymp}.(S_1 \sqcap \dots \sqcap S_5) \quad (6)$$

Using either (2–4), (6) or (2), (5–6) as the KB, we equally derive the hypothesis H_1 for both observations O_1 and O_2 , exactly as before.

As both of the above outlined solutions allow to derive the desired hypothesis, we will favour the solution on the right hand side, as it requires a less expressive DL (i.e., role transitivity axioms of the form (4) are not needed).

3.2 Hierarchy of Diagnoses

Potential diagnoses of our interest are listed in Table 1. These concepts can be readily taken from a number of medical ontologies. We chose to root them in SNOMED CT, where they are present as disorders. As in many cases in the medical domain, the main diagnosis (DM), has some subtypes which we want to distinguish. There are some alternate diagnoses (e.g., diabetes insipidus) which need to be rooted out or confirmed during the diagnostic process. And there are some related diagnoses (e.g., obesity, ketoacidosis) which may take part in the diagnosis of DM as relevant symptoms. The diagnoses thus form a hierarchy.

We will narrow down our focus on the diagnoses listed in Table 1. The respective part of the hierarchy is formalized using the following DL axioms:

$$DM1 \sqcup DM2 \sqcup GD \sqsubseteq DM \quad (7)$$

$$LADA \sqsubseteq DM1 \quad (8)$$

The hierarchy of diagnoses has significant influence on the generated hypotheses. For instance, considering the axioms (5–6) formalizing the symptoms of DM, together

Table 1. Relevant diagnoses in the use case

| Disorder (SNOMED) | DL |
|--|------|
| Diabetes mellitus | DM |
| Diabetes insipidus | D1 |
| Diabetes mellitus type 1 | DM1 |
| Diabetes mellitus type 2 | DM2 |
| Latent autoimmune diabetes mellitus in adult | LADA |
| Gestational diabetes | GD |
| Obesity | Ob |
| Ketoacidosis | KA |

with (7–8), if some of the symptoms of DM are observed (e.g., O_1 or O_2), the abductive reasoner will generate a number of diagnoses: $H_1 = \{p: \exists \text{hasDiag.DM}\}$ as before, but in addition also $H_2 = \{p: \exists \text{hasDiag.DM1}\}$, $H_3 = \{p: \exists \text{hasDiag.DM2}\}$, $H_4 = \{p: (\exists \text{hasDiag.DM1}) \sqcap (\exists \text{hasDiag.DM2})\}$, and similar hypotheses for all (asserted or derived) subconcepts of DM. This is because they now all allow to derive the given observations. On the other hand, if we compare these hypotheses semantically, we see that $H_i <_{\mathcal{K}} H_1$ for $i > 1$ in the hypotheses above, as $\mathcal{K} \cup H_i \models H_1$, and never vice versa. Hence only H_1 will be preferred.

3.3 Distinguishing Between Two Diagnoses

A typical task in medical diagnostics is to distinguish one diagnosis from another, often similar one. This is called differential diagnosis. The two diagnoses may be distinguished by considering some symptoms relevant to one of them, but not the other.

We will consider DM1 and DM2; as differentiating between them is a relevant medical problem. DM1 and DM2 have some symptoms in common, but they also have some different symptoms. In this case, the common symptoms are those of DM. When the patient has some of these symptoms we say that she has DM – this case is covered by axioms (5–6).

Specific symptoms of DM1 include: belly pain, vomiting, fruity breath odor, drowsiness, and coma (S_{10}, \dots, S_{14}). This is formalized as follows:

$$\exists \text{hasDiag.DM1} \sqsubseteq \exists \text{hasSymp.}(S_{10} \sqcap \dots \sqcap S_{14}) \quad (9)$$

Analogously, the symptoms of DM2 include: skin problems, slow healing, tingling, numbness, and high BMI. We will name these as S_{15}, \dots, S_{19} , which gives us the axiom:

$$\exists \text{hasDiag.DM2} \sqsubseteq \exists \text{hasSymp.}(S_{15} \sqcap \dots \sqcap S_{19}) \quad (10)$$

Together the three diagnoses are now covered with (5–7) and (9–10). Let us consider some observations and the respective hypotheses:

- If we observe some set of symptoms that are common to both DM1 and DM2, e.g., having again the observation $O_1 = p: (\exists \text{hasSymp.S}_2) \sqcap (\exists \text{hasSymp.S}_8)$, then the most preferred diagnosis will be $H_1 = \{p: \exists \text{hasDiag.DM}\}$. However, also $H_2 = \{p: \exists \text{hasDiag.DM1}\}$, $H_3 = \{p: \exists \text{hasDiag.DM2}\}$, $H_4 = \{p: (\exists \text{hasDiag.DM1}) \sqcap$

$(\exists \text{hasDiag.DM}_2)$, will be valid abductive explanations (among others), but as we already discussed in Sect. 3.2 they are all stronger than H_1 and hence H_1 will be the most preferred.

- If we observe at least one symptom specific to DM1, e.g., having the observation $O_3 = p: (\exists \text{hasSymp.S}_2) \sqcap (\exists \text{hasSymp.S}_{10})$, this is no longer abductively explained by H_1 , nor H_3 . The most preferred hypothesis will be H_2 . H_4 is an abductive explanation as well, but it is stronger than H_2 , and hence H_2 is preferred.
- The case when we observe some specific symptoms of DM2 (but none of DM1) is exactly analogous.
- Finally, if we observe specific symptoms of both DM1 and DM2, e.g., we have $O_4 = p: (\exists \text{hasSymp.S}_{10}) \sqcap (\exists \text{hasSymp.S}_{15})$, the most preferred hypothesis that abductively explains this observation is H_4 . This is because in case of H_1 – H_3 there is always some symptom which is not explained. There are other explanations (e.g., $H_5 = \{p: \exists \text{hasDiag.}(\text{DM1} \sqcap \text{DM2})\}$), but they are all stronger than H_4 .

3.4 Case of Secondary Explanation of the Observation

During differential diagnosis it is also important to recognize cases when the given set of observed symptoms may have other possible explanations than the disease in question. For instance, one of the major symptoms of DM (hyperglycemia named as S_3) may be caused as a side effect of some medication.

As the axiom (6) includes S_3 as one of the DM symptoms, we are already able to answer on the observation of having S_3 . Another possible explanation is taking the medications (name them M_1). So we add a new axiom:

$$\exists \text{hasMedication.M1} \sqsubseteq \exists \text{hasSymp.S}_3 \quad (11)$$

As this is a very simple conceptualization, there is only one observation by which this axiom play role. This observation is $O_5 = p: \exists \text{hasSymp.S}_3$. As well we have exactly two hypotheses $H_6 = \{p: \exists \text{hasMedication.M1}\}$ and $H_7 = \{p: \exists \text{hasDiag.DM}\}$. Neither H_6 nor H_7 is stronger then the other one so both are preferred.

This result means, that we cannot be sure, which explanation is correct and we have to continue in diagnosing. We are satisfied with this answer because also in medical domain information about hyperglycemia presence is insufficient condition to decide between taking medications and having diagnosis DM.

3.5 Associated Conditions

In the medical domain, relations between diagnoses may come into play. Some associated conditions may have similar symptoms, or subset of symptoms as other diagnoses. Thus, in fact, in Axiom (9) the symptoms S_{10}, \dots, S_{14} are symptoms of an associated diagnosis called ketoacidosis. Similarly the symptom S_{19} is a symptom of obesity, an associated diagnosis of DM2.

To take this into the account, we may add to (9) a new axiom, Axiom (12), and analogously to (10) a new axiom, Axiom (13), as follows:

$$\exists \text{hasDiag.KA} \sqsubseteq \exists \text{hasSymp.}(S_{10} \sqcap \dots \sqcap S_{14}) \quad (12)$$

$$\exists \text{hasDiag.Ob} \sqsubseteq \exists \text{hasSymp.S}_{19} \quad (13)$$

When we take the observations O_1 and O_3 and try to explain them using axioms (5–7) and (9–13), we see some changes:

- The case of observation $O_1 = p: (\exists \text{hasSymp.S}_2) \sqcap (\exists \text{hasSymp.S}_8)$ is not affected by the newly added axioms, as the symptom S_2 is not explained by any of them. So, the most preferred hypothesis is again H_1 .
- If we observe at least one symptom specific to ketoacidosis together with some symptoms of DM (which are all symptoms of DM1 as we already know), e.g., having the observation $O_3 = p: (\exists \text{hasSymp.S}_2) \sqcap (\exists \text{hasSymp.S}_{10})$, then besides for $H_2 = \{p: \exists \text{hasDiag.DM1}\}$ we have to consider also hypothesis $H_6 = \{p: (\exists \text{hasDiag.DM}) \sqcap (\exists \text{hasDiag.KA})\}$: they both explain O_3 and both are preferred to any other but they are mutually incomparable. But H_6 is certainly unexpected as so far we modelled the axioms in such a way that abductively the symptoms DM and ketoacidosis explain the diagnosis of DM1. To solve this we have to add yet another axiom:

$$\exists \text{hasDiag.DM} \sqcap \exists \text{hasDiag.KA} \sqsubseteq \exists \text{hasDiag.DM1} \quad (14)$$

The axiom enables to compare the two hypotheses also deductively, i.e., any patient with both diagnoses DM and ketoacidosis is inferred to have also DM1. Hence H_6 is now stronger than H_2 , and so H_2 is the single most preferred hypothesis.

- The case when we observe some specific symptoms of obesity together with symptoms of DM is analogous. We get the expected hypothesis $H_3 = \{p: \exists \text{hasDiag.DM2}\}$, but to suppress $H_7 = \{p: (\exists \text{hasDiag.DM}) \sqcap (\exists \text{hasDiag.Ob})\}$ as less preferred, we need to add:

$$\exists \text{hasDiag.DM} \sqcap \exists \text{hasDiag.Ob} \sqsubseteq \exists \text{hasDiag.DM2} \quad (15)$$

- If we observe only the symptoms shared by obesity and DM2, in our simplified example, only $O_5 = p: (\exists \text{hasSymp.S}_{18})$, two preferred and incomparable hypotheses are $H_8 = \{p: (\exists \text{hasDiag.Ob})\}$ and $H_3 = \{p: (\exists \text{hasDiag.DM2})\}$. In this case, this is in accord with the medical knowledge: this symptom can either be caused by one condition or by the other.

3.6 Case of Complications

However, in certain circumstances, one may wish to model the associated diagnoses differently, to capture a closer relation between them. This is, for instance, the case of DM1 and ketoacidosis. When we consider $O_3 = p: (\exists \text{hasSymp.S}_2) \sqcap (\exists \text{hasSymp.S}_{10})$ the modelling from the previous section gives the single most preferred diagnosis $H_2 = \{p: \exists \text{hasDiag.DM1}\}$. While ketoacidosis is also indicated by symptom S_{10} , the hypothesis $H_9 = \{p: \exists \text{hasDiag.KA}\}$ is not an abductive explanation as it does not explain S_{10} .

This solution does not correctly capture the importance of ketoacidosis presence in diabetic patients. Ketoacidosis does not merely share symptoms with DM1, but it is an acute complication thereof which rarely occurs in non-diabetic individuals.

Therefore, we would intuitively want the explanation $H_{10} = \{p: \exists \text{hasDiag.DM1}\} \sqcap (\exists \text{hasDiag.KA})$ in this case. In fact, H_{10} also explains O_3 in this case but is stronger than H_2 and hence also less preferred.

To achieve this, we have to remodel the knowledge, so that, given some symptoms that are shared by both DM1 and ketoacidosis, only the conjunction of these diagnoses explains them: we exchange (9) together with (12) with a new axiom:

$$(\exists \text{hasDiag.DM1}) \sqcap (\exists \text{hasDiag.KA}) \sqsubseteq \exists \text{hasSymptom.}(S_{10} \sqcap \dots \sqcap S_{14}) \quad (16)$$

This however still does not force H_{10} as single most preferred, as due to the presence (7) and (14) we now have $H_{10} <_{\mathcal{K}} H_6$ and $H_6 <_{\mathcal{K}} H_{10}$, hence both H_{10} and H_6 are equally preferred. However, recall that we have asserted (14) only to support a slightly different relation between DM1 and ketoacidosis in the previous section, so when we drop this axiom then the single, most preferred diagnosis remains to be H_{10} .

This last move reflects the fact that in order to model different relationships between preferred explanations we need to manipulate also the deductive part of the KB (used during deductive comparisons of hypotheses) and, what is more, these manipulation is not just additive, it is selective, which may cause problems in more complex cases with higher number of variously interrelated diagnoses.

3.7 Case of Further Examination Needed

Consider again the case of one hypothesis being a more specific case of another. Typically, we have two diagnoses which have some common symptoms, while the more specific one likely has some additional symptoms (like with DM1 and DM, or DM2 and DM above). If we observe only some of the common symptoms, our previous modelling derives the less specific hypothesis. The more specific hypothesis also explains the observations, but there is no additional evidence for it, so we prefer the less specific one.

Most of the time this is expected, but not all the time. It may be necessary to differentiate between such hypotheses by all means, and so, if perhaps some evidence is lacking it should be obtained by additional examination if possible, by a laboratory test for instance. Hence the outcome from the diagnosis procedure should not be just the less specific hypothesis, but instead it should indicate also that additional tests are needed.

From a medical perspective this might be illustrated on the following problem: some patients who fit a certain profile (they are older, and not obese) and typically show symptoms common to DM1, may in fact have a specific type of DM1 called LADA (cf. Table 1 and axiom (8)). To differentiate between these two diagnoses, medical practitioners are advised to test antibodies (e.g., GADA) in a blood sample.

As this case leads to some complex modelling, we will explain it on simplified examples. Assume that LADA is a specific form of DM1, and that DM1 has some symptom (S'_1) and LADA has an additional specific one (S'_2). That is, we start from:

$$\text{LADA} \sqsubseteq \text{DM1} \quad (17)$$

$$\exists \text{hasDiag.DM1} \sqsubseteq \exists \text{hasSymp.S}'_1 \quad (18)$$

$$\exists \text{hasDiag.LADA} \sqsubseteq \exists \text{hasSymp.S}'_2 \quad (19)$$

Now we end up exactly in the situation described above. If only DM1 symptoms are observed (i.e., in our simplified case $O_6 = p: \exists \text{hasSymp.S}'_1$) then the preferred hypothesis is $H_2 = \{p: \exists \text{hasDiag.DM1}\}$. In order to resolve this we may try to use (20) instead of (18):

$$(\exists \text{hasDiag.DM1}) \sqcap (\exists \text{needS.LT}) \sqsubseteq \exists \text{hasSymp.S}'_1 \quad (20)$$

We now get the expected most preferred hypothesis $H_{11} = \{p: (\exists \text{hasDiag.DM1}) \sqcap (\exists \text{needS.LT})\}$ for O_6 and so much for $O_7 = p: \exists \text{hasSymp.S}'_2$ we will get $H_{12} = \{p: \exists \text{hasDiag.LADA}\}$ as most preferred, however for $O_8 = p: (\exists \text{hasSymp.S}'_1) \sqcap (\exists \text{hasSymp.S}'_2)$ we now get $H_{12} = \{p: (\exists \text{hasDiag.LADA}) \sqcap (\exists \text{needS.LT})\}$ as most preferred, as the need of the lab test is now necessary to explain S'_1 . This is unintuitive as indeed once we observe S'_2 we know that the patient has LADA, no more tests are needed. It is easy to verify that it does not help to alter (19) to (21):

$$\exists \text{hasDiag.LADA} \sqsubseteq \exists \text{hasSymp.}(S'_1 \sqcap S'_2) \quad (21)$$

This is due to (17), (20), and (21) now for O_6 give both H_{11} and H_{12} as preferred, and they are incomparable. But clearly H_{12} is exactly the wrong hypothesis here. We are getting into some vicious circle.

The only solution we were able to come up with is to start treating the symptoms as completely specified. That is, either S'_1 or $\neg S'_1$ is always part of the observation, and same for S'_2 or other symptoms possibly involved in this part of the derivation. Using (17), (22), and (23), we always get the expected results:

$$(\exists \text{hasDiag.DM1}) \sqcap (\exists \text{needS.LT}) \sqsubseteq \exists \text{hasSymp.}(S'_1 \sqcap \neg S'_2) \quad (22)$$

$$\exists \text{hasDiag.LADA} \sqsubseteq \exists \text{hasSymp.}(S'_1 \sqcap S'_2) \sqcap \exists \text{hasSymp.}(\neg S'_1 \sqcap S'_2) \quad (23)$$

However, we now have to ask queries differently. For $O'_6 = p: \exists \text{hasSymp.}(S'_1 \sqcap \neg S'_2)$ the most preferred hypothesis is H_{11} , and for both for $O'_7 = p: \exists \text{hasSymp.}(\neg S'_1 \sqcap S'_2)$ and for O_8 (no need to change here) we will get H_{12} .

So, we were able to get the expected results, but for a considerable price. Treating symptoms as completely specified is not in line with the usual intuitions behind abduction, where it is normal to assume that the observations are incomplete, and we are tasked to give the most appropriate explanation for any such given observation. In addition, it leads to a considerable blow up in the axioms, where all possible combinations of positive and negative symptoms need to be enumerated.

4 Discussion

All formalizations in our use case are instances of the ABox abduction problem as defined in Definition 1, based on Elsenbroich et al. [7]. The explanations that we seek are typically of a specific form of a single assertion:

$$p: (\exists R_1.C_1) \sqcap \dots \sqcap (\exists R_n.C_n) \quad (24)$$

involving single patient p , where R_1, \dots, R_n most typically will come from some a priori known set of roles which are relevant based on the domain knowledge. A similar assumption may often hold also for the concepts C_1, \dots, C_n (e.g., most often these will be atomic concepts for various diagnoses and conditions relevant to the patients state).

While exceptions to these constraints may certainly be found (e.g., a second person involved from which the patient contracted the disease), in many cases the form of expected explanations will be reducible into atomic diagnoses by adding “interface” axioms of the form:

$$\text{Diag1} \equiv (\exists R_1.C_1) \sqcap \dots \sqcap (\exists R_n.C_n) \quad (25)$$

The hypothesis of the complex form (24) now reduces into the atomic form p : Diag1 which makes it possible to postulate the problem as simple ABox abduction. This is an important observation, as Du et al. [5] showed that the simple abduction problem can be effectively solved even for DLs up to *SHIQ*.

Most of the examples we have shown rely in a fairly basic DL *ALC*, abduction support for which is known [12,13,5,9]. We used complex role inclusions to capture dependencies between symptoms, but we also showed a simpler modelling which does not require it. Different complex DL constructs that might possibly be needed in more realistic situations certainly include number restrictions, and also restriction over concrete domains (known, e.g., in OWL [14]), that would enable to support some more elaborate statements about symptoms (the patient has at least some number of symptoms, or has a numeric value of some symptom from a specific range). Extensions of abductive reasoning for more expressive DLs that include such constructs are therefore desirable.

From a modelling perspective, an interesting lesson learned from the use case is that modelling a knowledge base to be used in a classical, deductive way, and modelling it to support abductive reasoning pose different, and sometimes conflicting requirements. As we noted in Sect. 3.1, it is often the case that certain explanation is expected to be valid for a number of symptoms, including any subsets thereof. Hence the typical approach that we relied upon is to formulate the axioms in a stronger fashion – the explanation implies all the symptoms, hence any subset follows. This kind of modelling is also demonstrated in the literature [7,5].

Firstly, we note that this requires the knowledge to be used in abductive application to be modelled differently than it is typically usual for ontologies, which normally are subject to deductive reasoning. Secondly, as we have repeatedly observed in this paper, even the process of abduction has an important subtask when hypotheses are compared and strictly deductive reasoning is used for this. This results in complex inter-relations between the “abductive” and the “deductive” part of the knowledge base and may lead to rather complicated modelling. A possible way how to deal with this is to treat the abductive and the deductive knowledge separately [16], or to use a more refined formulation of the abduction problem [10]. We plan to try this in the future.

Finally, we note that abduction has also a number of advantages, e.g., when compared to deductive diagnostic reasoning. In number of works [8,17,18] relying upon the latter, where deductive rules of the form $S_1, \dots, S_n \rightarrow D$ are used, where S_i are symptoms and D is the diagnosis, the authors point out the problem occurring in case of two diagnoses D_1, D_2 with the former having a subset of symptoms of the latter. Using deductive inference, and observing all symptoms of D_2 , both D_1 and D_2 are derived. This is counterintuitive as specific symptoms of D_2 were observed which are not symptoms of D_1 . This problem has to be addressed by some suitable workarounds. In comparison, as we have demonstrated in the use case, abductive reasoning naturally eliminates the hypotheses which do not explain all of the observed symptoms.

Acknowledgments. This work was supported by VEGA project no. 1/1333/12. Júlia Pukancová is also supported by grant GUK no. UK/426/2015.

References

1. American Association of Clinical Endocrinologists: Diabetes care plan guidelines. *Endocrine Practice* 17 (2011)
2. American diabetes association: Diagnosis and classification of diabetes mellitus. *Diabetes Care* 31 (2008)
3. Baader, F., Calvanese, D., McGuinness, D.L., Nardi, D., Patel-Schneider, P.F. (eds.): *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press (2003)
4. Colucci, S., Noia, T.D., Sciascio, E.D., Donini, F.M., Mongiello, M.: Concept abduction and contraction in description logics. In: *Proceedings of the 2003 International Workshop on Description Logics (DL2003)*, Rome, Italy September 5-7, 2003 (2003)
5. Du, J., Qi, G., Shen, Y., Pan, J.Z.: Towards practical ABox abduction in large description logic ontologies. *Int. J. Semantic Web Inf. Syst.* 8(2), 1–33 (2012)
6. Eiter, T., Gottlob, G., Leone, N.: Abduction from logic programs: Semantics and complexity. *Theor. Comput. Sci.* 189(1-2), 129–177 (1997)
7. Elsenbroich, C., Kutz, O., Sattler, U.: A case for abductive reasoning over ontologies. In: *Proceedings of the OWLED*06 Workshop on OWL: Experiences and Directions*, Athens, Georgia, USA, November 10-11, 2006 (2006)
8. García-Crespo, Á., Rodríguez, A., Mencke, M., Gómez-Berbís, J.M., Colomo-Palacios, R.: Oddin: Ontology-driven differential diagnosis based on logical inference and probabilistic refinements. *Expert Systems with Applications* 37(3), 2621–2628 (2010)
9. Halland, K., Britz, K.: Naïve ABox abduction in \mathcal{ALC} using a DL tableau. In: *Proceedings of the 2012 International Workshop on Description Logics, DL-2012*, Rome, Italy, June 7-10, 2012. Sun SITE Central Europe (CEUR) (2012)
10. Hubauer, T., Lamparter, S., Pirker, M.: Relaxed abduction: Robust information interpretation for incomplete models. In: *Proceedings of the 24th International Workshop on Description Logics (DL 2011)*, Barcelona, Spain, July 13-16, 2011 (2011)
11. Hubauer, T., Legat, C., Seitz, C.: Empowering adaptive manufacturing with interactive diagnostics: A multi-agent approach. In: *Advances on Practical Applications of Agents and Multiagent Systems - 9th International Conference on Practical Applications of Agents and Multiagent Systems, PAAMS 2011*, Salamanca, Spain, 6-8 April 2011. pp. 47–56 (2011)
12. Klarman, S., Endriss, U., Schlobach, S.: ABox abduction in the description logic \mathcal{ALC} . *Journal of Automated Reasoning* 46(1), 43–80 (2011)
13. Ma, Y., Gu, T., Xu, B., Chang, L.: An ABox abduction algorithm for the description logic \mathcal{ALCI} . In: *Intelligent Information Processing VI - 7th IFIP TC 12 International Conference*. pp. 125–130 (2012)
14. OWL Working Group (ed.): *OWL 2 Web Ontology Language Document Overview*. Recommendation, W3C (27 October 2009)
15. Peirce, C.S.: Deduction, induction, and hypothesis. *Popular science monthly* 13, 470–482 (1878)
16. Petasis, G., Möller, R., Karkaletsis, V.: BOEMIE: Reasoning-based information extraction. In: *Proceedings of the 1st Workshop on Natural Language Processing and Automated Reasoning co-located with 12th International Conference on Logic Programming and Non-monotonic Reasoning (LPNMR 2013)*, A Corunna, Spain, September 15th, 2013. pp. 60–75 (2013)
17. Rodríguez, A., Labra, J., Alor-Hernandez, G., Gómez, J.M., Posada-Gomez, R.: ADONIS: Automated diagnosis system based on sound and precise logical descriptions. In: *Proc. of 22nd IEEE International Symposium* (2009)

18. Rodríguez-González, A., Labra-Gayo, J.E., Colomo-Palacios, R., Mayer, M.A., Gómez-Berbís, J.M., García-Crespo, A.: SeDeLo: Using semantics and description logics to support aided clinical diagnosis. *Journal of Medical Systems* 36(4), 2471–2481 (2012)