PLAUSIBLE REASONING OVER KNOWLEDGE GRAPHS: A NOVEL APPROACH FOR SEMANTICS-BASED HEALTH DATA ANLYTICS

by

Hossein Mohammadhassanzadeh

Submitted in partial fulfilment of the requirements
for the degree of Doctor of Philosophy

at

Dalhousie University
Halifax, Nova Scotia
March 2019

To Maman and Baba
for their love, encouragement
and unwavering support,

To my siblings
for helping me dream big,

To my grandparents
for teaching me hard work.

Table of Contents

List of Tables

List of Figures

Abstract

The Semantic Web is regarded as the next generation of World Wide Web, in which human and machine readable and understandable knowledge is exchanged. The Semantic Web allows the generation of new knowledge by analyzing the underlying semantics through a variety of reasoning mechanisms. Plausible reasoning provides a non-deductive and exploratory approach to infer new knowledge from large data-sets. Plausible reasoning generates meaningful associations between data elements by analyzing the semantics of the data to identify plausible knowledge that can assist in complex decision making, especially when dealing with incomplete knowledge. Hence, plausible reasoning is an interesting and viable approach for semantic data analytics, providing an exploratory approach to 'Big' data analytics.

In this thesis, we investigate plausible reasoning for semantic data analytics focusing on (a) identification and formal definition of plausible patterns; (b) implementation of a plausible reasoning framework capable of providing semantic analytics; and (c) evaluation of the efficacy of plausible reasoning for analyzing large volumes of health data.

We used knowledge graphs, a Semantic Web inspired knowledge representation formalism, to encode semantic associations between entities. To infer new knowledge, we identified six plausible patterns—i.e. generalization, specialization, interpolation, a fortiori, (dis)similarity, that are applied to three types of semantic relationships—i.e. conceptual hierarchy, partial order and equivalence. We developed a plausible extension to the Web Ontology Language (OWL) in terms of PL-OWL to represent order-based relationships. The plausible patterns are employed by our SeDan (SEmantics-based Data Analytics) framework that uses the OWL 2 QL profile (underpinned by DL-Lite family) to support query answering over knowledge graphs.

To evaluate our approach, we designed a real-world medical setting in which SeDan is required to answer intelligent medical questions from BioASQ challenges, using the large-scale SemanticMEDLINE database, while the standard clinical ontologies, DrugBank and Disease Ontology, provide the supplementary semantics. In addition to providing plausibly inferred answers, the correctness of the answers and the underlying reasoning processes are important. The experimental results show SeDan expands the query answering coverage of the database by 37 percent, while 88 percent of the plausible answers are clinically reasonable, verified by a domain expert.

List of Abbreviations Used

DAG	Directed Acyclic Graph
DL	Description Logic
DQ	Datalog Query
EHR	Electronic Health Record
EMR	Electronic Medical Record
KBS	Knowledge Based System
KG	Knowledge Graph
LUBM	Lehigh University Benchmark
OBDA	Ontology-Based Data Access
OWA	Open World Assumption
OWL	Web Ontology Language
PL	Plausible
PL-OWL	Plausible OWL
PL-QR	Plausible Query Rewriting
PR	Plausible Reasoning
PR	Plausible Reasoning
QR	Query Rewriting
RDBMS	Relational Data Base Management System
RDF	Resource Description Framework
RDFS	RDF Schema
RuleML	Rule Markup Language
SeDan	SEmantics-based Data ANalytics
SPARQL	SPARQL Protocol and RDF Query Language
SPIN	SPARQL Inferencing Notation
SW	Semantic Web
SWRL	Semantic Web Rule Language
UCQ	Union of Conjunctive Query
URI	Uniform Resource Identifier
W3C	World Wide Web Consortium

Acknowledgements

I would never have been able to finish this journey without the support of a few special people. I hope the completion of this thesis compensates your effort.

First and foremost, I would like to express my deepest gratitude to my supervisor, Dr. Raza Abidi, for his excellent guidance, support, patience and enthusiasm. I am forever grateful for the opportunity you have provided me and the lessons I learned from you.

My sincere thanks also go to Dr. Samina Abidi for her valuable guidance, insightful comments and encouragement. The completion of this work would not have been possible without you imparting your knowledge and expertise in this study.

I would like to extend my appreciation to the rest of my supervisory committee, Dr. Malcolm Heywood, Dr. Christian Blouin, and the past members of the committee Dr. Mike Shepherd, and Dr. Denis Riordan, and as well, Dr. Wojtek Michalowski for kindly accepting to be my thesis external examiner. I thank you all for your contributions and valuable suggestions.

I am fully indebted to two colleagues, William Van Woensel and Enayat Rajabi, who challenged me and brought a depth of knowledge to this dissertation.

A very special thanks to my parents, Soudabeh and Khosro, who made all this possible. Your unconditional love, spiritual support, endless care and trust have made me a stronger person. I cannot thank you enough. I am also grateful to my siblings for always being there. Anything good that has come to my life has been, undoubtedly, because of my family and my grandparents' prayers. I hope I have made you all proud.

I also want to thank my friends, especially Farzaneh, Hossein, Ehsan N., Ehsan D., the Lahsaees, Mehrdad and Paige for their support along the way. I have been fortunate to have you as my friends.

Chapter 1: Introduction

Data analysis includes the interpretation of data aiming to derive new insights and induce further interesting knowledge. Exploratory Data Analytics (EDA) approaches, like plausible reasoning, are akin to natural experiments, in which the investigator leverages the available data to make assumptions, examine possible outcomes and present alternatives. The gained insights may be plausible, but they complement hypothetico-deductive approaches by generating data-driven inferences that identify underlying patterns and correlations between data elements that can be further used for analytical query answering over large-scale data (Ge, 2017; Kell & Oliver, 2004; Panahiazar, Taslimitehrani, Jadhav, & Pathak, 2014; Tukey, 1977).

Exploratory data analytics approaches combine data sources with data-driven methods of thinking, reasoning and analysis, which are different from traditional statistics and hypothesis testing techniques (Krumholz, 2014). Effective exploratory data analytics leverages the semantics of the concepts and their relationships, as represented in the data, to explore the underlying data (Tickoo & Iyer, 2017). This in-depth data analysis approach, not only enriches the description and interpretation of data, but more importantly offers the ability to derive inferences about the nature of the data using a range of logical reasoning approaches (Ogiela, 2013).

Plausible reasoning, as an exploratory data analytics approach, pertains to human's problem-solving process which explores associations between the underlying domain specific data in conjunction with background domain knowledge to discover 'plausible' inferences. Plausible reasoning explores a (partial) set of true statements to discover a *plausibly* acceptable association, which is the best-effort answer and often reasonable, considering the current understanding of the domain and the problem being investigated. In the absence of deterministic and 'complete' knowledge, plausibility can be regarded as naturally-fit solutions that can be deemed to be pragmatic and reasonable to solve a problem.

In functional terms, plausible reasoning derives plausible solutions by exploring semantic associations between the data—this is different from identifying frequent patterns or associations within the data—and using these semantic associations to derive inferences. Plausible reasoning, therefore, relies on domain-specific (conceptual) relational knowledge (e.g., conceptual hierarchy, partial order and equivalence) and that is used to identify a set of plausible patterns (such as a fortiori) inherent within the data. Although, plausible solutions may not always be supported by objective facts, their presence provides a way forward to solve complex problems (Habicht, Victora, & Vaughan, 1999).

Medical big data analysis can provide meaningful insights by turning collected data into actionable knowledge. The acquired knowledge helps to better informed clinical diagnoses, improve targeted (personalized) therapies, validate medical treatment and predict the adverse events to treatments, while lowering costs (Panahiazar et al., 2014; Roski, Bo-Linn, & Andrews, 2014; Weil, 2014). The recent surge in P4-medicine (Predictive, Preventive, Participatory, Personalized), the exploitation of smart devices and the ease of electronic communication have led to generation of large volumes of health-related data, which is by nature high in variety and velocity (Hood & Friend, 2011; Roski et al., 2014). As opposed to the challenges with the management, storage and processing of patient data, the large volume of data offers unprecedented opportunities to discover new relationships between data elements that can help advance our understanding of biological structures and clinical processes. (S. Abidi, Vallis, Piccinini-Vallis, Imran, & Abidi, 2018; Mathew & Pillai, 2016; Panahiazar et al., 2014).

Given the scope of medical knowledge, the potential to derive new knowledge using data-driven approaches is not just interesting but extremely useful (S. S. R. Abidi, 2007). Conventional health data analysis methods are hypothesis-driven approaches based on deductive reasoning conducted on relatively small amounts of data, drawing logical assertions that may have already been embedded in the premises (Morgenthaler, 2009; Roski et al., 2014). However, with the availability of large volumes of health data about clinical practices and processes, there is a need to set-aside traditional mind-sets and investigate additional reasoning methods that can infer plausible solutions. Whereas the truth of these solutions is not fully verified, their value is nevertheless based on observations from actual clinical practices that resulted in positive health outcomes (Krumholz, 2014).

Knowledge Graphs (KG) are an upcoming approach to represent massive volumes of semantic data by encoding conceptual entities, their properties and the chain of relationships connecting them. The connectivity of knowledge graphs offers the opportunity to identify interesting and unknown connections among data elements. The Semantic Web framework offers knowledge representation formalism, such as ontology languages with different level of expressivity, e.g., Resource Description Framework (RDF), RDF Schema and Web Ontology Language (OWL)—that are essential for semantic analytics via plausible reasoning. Furthermore, semantic web offers a range of reasoning mechanisms to perform reasoning over knowledge graphs, and in turn to implement plausible reasoning (Haider, Abidi, Woensel, & Abidi, 2014).

In this thesis, we investigate the potential of implementing plausible reasoning, as a viable exploratory reasoning method, to analyze semantically annotated (large) data (i.e., represented in knowledge graphs) to infer new knowledge, deliver useful insights and solve complex problems. We propose the concept of *semantics-based data analytics* seeking actionable insight from large data jointly with background domain knowledge, when available. In line with this objective, we evaluate the ability of our implementation of plausible reasoning using health data, aiming to transform large amounts of health data into insightful actions that can assist healthcare providers with better disease diagnosis and long-term care.

1.1 Research challenges

To accomplish the objective of this thesis, we have taken a Semantic Web inspired knowledge management approach to pursue our investigation of plausible reasoning over knowledge graphs. In this regard, there are few key research challenges being pursued:

a) The formal description of plausible reasoning; Plausible reasoning is an old reasoning approach widely used in different domains (e.g., philosophy, law, mathematics and artificial intelligence), while its definition has not always been unambiguous and consistently agreed upon. Thus, plausible reasoning and its characteristics that distinguish it from other sort of non-deductive reasonings should be properly studied and a formal model of plausible reasoning applicable to computer systems must be defined.

b) Identification of plausible patterns and their representation; Plausible reasoning performs inference through a set of frequently recurring patterns to suggest a plausible statement. There exists a variety of plausible patterns that not all of them are applicable to our purpose or feasible to formalize (i.e., in computer systems). Hence, identification of the various plausible patterns along with their functionality and the semantics that they are applied to are paramount.

c) Implementation of a purposeful plausible reasoning framework capable of providing semantic analytics; An effective development of plausible reasoning depends on (i) full support of the semantic associations exploited in the plausible patterns in one coherent framework, and (ii) formalizing and implementing the logic of plausible reasoning (i.e., analysis of semantically annotated data to generate new knowledge) exploiting reasoning methods offered by the Semantic Web. And,

d) Evaluating the efficacy of plausible reasoning for analyzing large volumes of health data for discovery of causal relationships between data elements (e.g., drug-disease causal relationship).

1.2 Contributions

Along with the primary contribution of this thesis on implementing plausible reasoning over knowledge graphs, this work also offers:

1. PL-OWL; a *plausible extension to OWL* ontology to support the full representation of semantic associations applicable to plausible patterns. An ontology language (i.e., OWL) should provide enough expressivity to simultaneously represent

various semantics (e.g., hierarchical, equivalence and ordered-based) applicable to plausible patterns in one integrated language.

2. A novel query-rewriting algorithm that implements plausible patterns within the Semantic Web framework using OWL 2 QL together with our PL-OWL. The *plausible query rewriting algorithm* reformulates SPARQL queries to explore knowledge graphs, discover logical entailments and draw plausible associations.
3. Design and development of a *SEmantics-based Data ANalytics (SeDan) framework* that integrates plausible reasoning with Semantic Web technologies to discover hidden knowledge underlying domain-specific data. SeDan performs an exploratory analysis on knowledge graphs to infer plausible knowledge and, ultimately, extend the query answering capabilities over knowledge graphs.

1.3 Thesis organization

The rest of this thesis proceeds as follows. First, Chapter 2 provides an overview of some basic concepts justifying the motivation of the thesis. The Semantic Web and its potential benefits to exploratory data analytics over knowledge graphs is discussed in Chapter 2. Chapter 3 introduces the notion of plausibility and focuses on the definition, and characterization of plausible reasoning and its components, trying to depict a clear and unambiguous picture of plausible reasoning. Having all the basic concepts introduced, Chapter 4 further elaborates the challenges of implementing plausible reasoning over knowledge graphs and presents our solutions addressing those challenges: the plausible extension to OWL and plausible query rewriting algorithm. Furthermore, the architecture of SeDan, our framework to semantics-based data analytics, is presented in Chapter 4. The evaluation framework of the system, the design of the experiment and the required materials are explained in Chapter 5. Also, this chapter provides the experimental results and discusses the findings. Conclusions and potentially useful future work are discussed in Chapter 6.

Chapter 2: Preliminaries

Given the subjective nature of plausible reasoning, its capability to derive reasonable inferences is constrained to realize which data is most relevant and accurate. To identify the relevancy of data, the meaning of data and its relationships with other data items are paramount. This chapter discusses how semantic associations capture the meaning of data and how the analysis of semantic data, especially within large-scale datasets represented as knowledge graphs, provides the opportunity to discover valuable and actionable insights, and how the Semantic Web technologies contribute to this in-depth analysis.

This chapter introduces the basic concepts that this dissertation is built upon and later in the following chapters we show how these building blocks justify the motivation and contributions of the work. However, we cannot study the analysis of data without considering that data is inherently uncertain and incomplete. The challenges of reasoning with uncertainty and the approaches to address it are discussed in this section as well.

2.1 Semantic association

Semantics denote the meaning of data, rather than its syntax or structure. Semantic associations imply (complex) relationships between resource entities. In particular, semantic associations express meaningful relationships between two or multiple concepts, which can be represented as a directed labeled link between the concepts/resources in a semantic space (i.e., in the form of concept-relation-concept triple). In general, a semantic association can be expressed as a meaningful path between any two entities or resources (Anyanwu & Sheth, 2003; Khoo & Na, 2006; Kim, Ostrowski, Yamaguchi, & Sheu, 2013).

In a graph representation of knowledge (e.g., RDF graph), two entities, ν_1 and ν_n, are semantically associated, $\rho(\nu_1, \nu_n)$, if there exists a sequence of relationships σ_i $(1 \leq i)$ which $\nu_1, \sigma_1, \nu_2, \sigma_2, \nu_3, \sigma_3, \ldots, \nu_{n-1}, \sigma_{n-1}, \nu_n$. In a knowledge graph, a semantic association may imply an actionable insight representing the connections between different events, concepts or phenomena, which also can be inferred either based on data or domain knowledge (B. Aleman-Meza, Halaschek-Wiener, Arpinar, Ramakrishnan, & Sheth, 2005; Boanerges Aleman-Meza, Sheth, Palaniswami, Eavenson, & Arpinar, 2006; Matthew Perry, Sheth, Arpinar, & Hakimpour, 2009).

2.2 Semantic analytics

Analysis of data at the semantic level will provide new opportunities to generate and discover information and transform information into actionable knowledge (Serrano & Gyrard, 2016). Semantic analytics performs information analysis by investigating the relationships between different entities in ontologies and semantic metadata. Analysis of semantic data exploits named relationships with well-defined semantics, which makes it distinguishable from statistical approaches of data mining and machine learning (Matthew Perry et al., 2009), and typical querying and inferencing mechanisms (Decker, 2007; Serrano & Gyrard, 2016).

In the analysis of semantic data, the representation of associations (i.e., complex relationship between two entities) is a key issue. Semantic Web knowledge representation models (e.g., such as RDF(S) and OWL) offer efficient tools to representation and analysis of semantically annotated data (e.g., RDF data), seeking actionable knowledge. In these languages data associations are the fundamental elements making the query and analysis of data easy, while keeping it understandable for humans (Matthew Perry et al., 2009).

2.2.1 Related work

Mehdi et. al. (Mehdi, Brandt, Roshchin, & Runkler, 2016) investigated the potential of semantic technology to interact and leverage data analytics for operational use. In their research, they aimed to exploit massive industrial data from sensors and devices (i.e., Siemens Turbine), providing insights into real-time system conditions, enhanced decision support, reliability and cost reduction. Mehdi et. al. combined data-driven strategies with knowledge models by introducing an upper-level ontology of a technical system, expressed in OWL 2 QL. In their semantic framework for turbine analytics, the semantic layer, in the middle of the framework, performs semantic mapping based on the introduced ontology. An Ontology-Based Data Access (OBDA) system (Kharlamov et al., 2013) will query the data based on the domain-specific language rather than the actual heterogeneous data sources. In their model, Mehdi et. al. leveraged Semantic Web technologies with specialized programs to reduce the complexity problems of big industrial data analytics. However, Mehdi et. al. just introduced a generic ontology to overcome the complexity problem, and no analysis of semantic data leading to new inferences is performed.

Zimmermann et. al. (Zimmermann, Lopes, Polleres, & Straccia, 2012) presented an extension of RDFS to support meta information (semantics) from three domains of temporal, fuzzy, provenance or any combination of them (e.g., temporally-annotated fuzzy) in the form of RDF annotations. For example, (s, p, o): $[\lambda]$ identifies a semantically annotated triple with meta information where $[\lambda]$ can be a *confidence* value or a *timestamp* representing a time point or a *temporal interval*. To be able to query and reason over the semantically annotated RDF, Zimmermann et. al. also introduced AnQL, as an extension to SPARQL. In their approach, the initial effort to enrich the triple statements in RDF repository with the semantic annotations is high. Also, more importantly, the considerations in writing the query using AnQL make the query generation complicated.

Likewise, there exist a number of frameworks offering semantic analytics, which exploit domain-specific ontologies to enrich search algorithms (Ślęzak et al., 2011) (e.g., search for a document), question answering (MS Perry, 2008) (e.g., answering biomedical questions), information retrieval (Martin, Weibel, Rocke, & Boker, 2018; Ramakrishnan, Kochut, & Sheth, 2006), and visual analytics (Endert, 2016; Hsiao, Pandhalkudi Govindarajan, & Lin, 2016).

The studies above, along with many other works, leverage semantically annotated data to facilitate the analysis of large-scale data. They demonstrate the utility of semantics in getting to know the data, providing different perspectives on data, and adding highly valuable information helping to better interpret underlying data when it comes to decision making and problem solving. However, there still exist a lack of mechanisms to in-depth exploratory data analysis, that not only enriches the description and interpretation of data, but more importantly offers the ability to derive inferences about the nature of the data and yield actionable descriptive, predictive and prescriptive solutions using a range of logical reasoning approaches (Kaisler, Espinosa, Armour, & Money, 2014; Ogiela, 2013). In this regard, an exploratory data-driven reasoning approach to big data analytics that allows the generation of new knowledge by analyzing the underlying semantics of the available data/knowledge is paramount.

2.3 Semantic Web

Semantic Web technologies offer languages that express semantic meta-data for human and machine consumption. Enriching a Web of resources with metadata, helps Semantic Web technologies to simultaneously carry syntactics and semantics of data. Tim Berners-Lee (Tim Berners-Lee, Hendler, & Lassila, 2001) define the Semantic Web as:

> *"The Semantic Web is an extension of the current Web in which information is given a well-defined meaning, better enabling computers and people to work in cooperation."*

Conforming with the Open World Assumption (OWA), Semantic Web (SW) provides an environment in which unknown facts are not considered false, but assumed not yet discovered (Sabou, 2016). With this assumption, the Semantic Web technologies have transformed the process of data acquisition, integration, inquiry and facilitates the development of open-minded knowledge discovery/reasoning/inference techniques required in the exploratory data analytics methods (Ge, 2017).

The World Wide Web Consortium (W3C) is the main organization developing standards for the World Wide Web including the Semantic Web. To deliver the Semantic Web vision, W3C introduces a series of languages, standards, recommendations, frameworks, and APIs.

Knowledge representation formalisms in the Semantic Web framework offer representations that concurrently capture syntax of data while carrying its semantics. The Resource Description Framework (RDF), RDF Schema (RDFS) and Web Ontology Language (OWL) represent fine-grained annotated data at various levels of expressivity.

The Resource Description Framework (RDF) (Consortium, 2014) is the basic data model representing both data and meta-data. RDF, which was basically an attempt to tackle semantic limitations of XML, represents information in a way that it is syntactically and semantically understandable, interoperable, and reusable both for humans and machines. RDF statements represent both the properties associated with the entities and their relationships with other entities (Heflin, 2001; Wu, Eadon, & Das, 2008). RDF repositories store data in triple format and construct a flexible and extensible knowledge representation formalism that facilitates the incorporation of newly discovered facts.

RDF Schema (RDFS) and OWL ontologies represent the associations/relationships between the concepts, while seamlessly capturing their semantics in an accurate, rich and unambiguous way. RDFS defines a basic *vocabulary* describing RDF data model. RDFS is a simple ontology language to describe expressive taxonomies, which is written in RDF, providing a mechanism to define groups of related concepts and their relationships (Abell et al., 2017; Chiba, Nishide, & Uchiyama, 2015). RDFS organizes the classes and properties in a hierarchical format. Class definitions include *Resource*, *Class*, *Literal*, *Datatype*, *Property*, etc. Properties, which are the instances of the class *rdfs:Property*, define a relationship between two class of concepts. Properties like *domain*, *range*, *subClassOf*, *subPropertyOf*, etc. are some property constructs of RDFS. The basic ontological constructs of RDFS limits its reasoning capabilities to very basic inferences about taxonomies (Horrocks, 2003).

Web Ontology Language (OWL) (Harmelen & McGuinness, 2004) is an Description Logic-based extension to RDF and RDFS offering flexible ontology modeling while providing efficient automated reasoning. Well-defined syntax and semantics of OWL is inherited from its predecessor DAML+OIL, providing an efficient reasoning with enough expressivity. Essentially, OWL uses RDF statements and supports all the class and property constructs of RDFS, while delivering a richer expressivity. OWL provides logical combinations of other classes (e.g., *intersections*, *unions*, or *complements*) and makes the enumeration of the objects feasible. In OWL, *equality* and *disjoint* statements can be made on classes and properties. The major addition of OWL is its capability to define restrictions on a particular property of a class to determine how properties behave. OWL has three sublanguages: Lite, DL, and Full, ordered in increasing expressivity. Because OWL Full is undecidable, the focus has been on OWL Lite and OWL DL which are based on expressive Description Logics (DL). Decidability and provision of sound, complete and tractable reasoning services are the key inference problems that motivate the use of different OWL sublanguages (Horrocks, 2003; Horrocks & Sattler, 2007).

Linked Open Data (LOD) methods have transformed data/knowledge sharing and are used to link diverse data sources and integrate background knowledge to further extend the data coverage. SPARQL query language is utilized to retrieve and manipulate the data stored in heterogenous, distributed triple repositories.

In addition, the knowledge representation formalisms of the Semantic Web allow automatic reasoning and inference over RDF triple stores (T Berners-Lee, Hendler, & Lassila, 2001; Laborie, Ravat, Song, & Teste, 2015). Built-in description logic-based reasoning, which supports OWL ontologies and Semantic Web rule languages (i.e., RuleML, SWIRL), helps perform semantic-aware analysis and automatically discovers relevant knowledge from underlying data/knowledge (Bouamrane, Rector, & Hurrell, 2011; Gnanambal, 2014; Mohammadhassanzadeh, Van Woensel, Abidi, & Abidi, 2017;

Rodríguez-González et al., 2012). The following sections further elaborate on the opportunities and challenges of reasoning on the Semantic Web.

2.3.1 Reasoning paradigms in the Semantic Web

Semantics and expressivity of RDFS and OWL constructs provide the opportunity of ontology-based reasoning to discover new relationships based on RDF data. This inference is a logical entailment of the semantic model (i.e., ontology) leveraging the stored data. The OWL constructs like *sameAs*, *differentFrom*, *disjointFrom*, *equivalentProperty*, *equivalentClass*, *FunctionalProperty*, *TransitiveProperty*, *someValuesFrom*, and RDFS vocabulary like *subClassOf*, *subPropertyOf*, *range*, *domain*, and more, are some constructs that deliver ontological inference on the Semantic Web.

OWL 2 EL, OWL 2 QL, and OWL 2 RL are three fragments (profiles) of OWL 2 DL introduced by W3C that target different application areas (i.e., regarding the expressivity of the ontology and size of the data) adjusting the expressivity power of OWL Full ontology to perform efficient reasoning (S. S. R. Abidi & Abidi, 2013; S. S. R. Abidi & Shayegani, 2009). OWL 2 EL performs reasoning in polynomial time suitable for applications with a large number of concepts and/or properties in their ontologies. OWL 2 RL is recommended for applications that require scalable and efficient reasoning by trading the full expressivity of the language (i.e., compared to OWL 2), while being amenable to implementation using rule-based technologies. OWL 2 RL runs efficient reasoning in polynomial time. The third fragment of OWL 2, OWL 2 QL, is designed as an ontology language that provides sound and complete query answering over (very) large data in LOGSPACE time (i.e., with respect to the size of the data). OWL 2 QL is based on DL-

Lite family providing many features required to query large data via a simple query rewriting mechanism (World Wide Web Consortium, 2012).

The choice of an OWL flavor strongly depends on the application, the size of the data and complexity of the ontology representing the domain knowledge: OWL 2 EL is suitable for complex ontologies with intricately related classes and properties; OWL 2 RL is recommended for applications which require additional expressivity in the form of rules; and if the main purpose of the application is to leverage a moderately complex ontology to reason over a large amount of instance data via query rewriting, OWL 2 QL is an appropriate choice.

Ontology languages concentrate on a formal specification of a conceptualisation. An ontology contains a set of class definitions, along with their attributes, associations between the classes (i.e., hierarchies) and the restrictions characterizing the relationships between classes and their instances (Lin, Harding, & Tsai, 2012). However, even the most expressive ontology language does not have enough capability to represent complex domain knowledge—e.g., OWL does not have the expressivity power to join relationships (like UncleOf relation). Ontology reasoning is based on description logics that are designed to acquire high expressiveness while maintaining the decidability of the reasoning. Classical ontology reasoning in the Semantic Web is limited to consistency checking, class properties and relationships and instance classification (Lin et al., 2012).

Integrating OWL ontology with rule-based representation of knowledge (i.e., in the form of implication and conjunction) can overcome this limitation (Jafarpour, Abidi, & Abidi, 2016; Matheus et al., 2005; Van Woensel, Roy, & Abidi, 2016). Rule-based reasoning offers more complex reasoning tasks than what ontology (e.g., OWL) reasonings deliver.

Rule languages allow users to describe relationships that cannot be described using DL used in OWL (Rattanasawad, Buranarach, Saikaew, & Supnithi, 2018). Semantic Web Rule Language (SWRL) supplements OWL profiles (i.e., OWL DL and OWL Lite profiles) with the Unary/Binary Datalog RuleML sublanguages of the RuleML providing a Horn clause rules extension to OWL. SWRL provides the required extra expressivity to reason with OWL individuals (Horrocks & Patel-Schneider, 2004).

In addition to ontology reasoning and rule-based reasoning, the latest version of SPARQL query language (SPARQL 1.1) offers *entailment regime*, which captures the capabilities of W3C standards layered on top of RDF (like OWL 2) in the process of query answering. This capability provides a flexible mechanism for extending SPARQL queries using ontologies that allow semantic interpretations of RDF graphs and ultimately allow inference of additional RDF statements from explicitly given assertions. One of the main motivations behind the entailment regime was to deal with the inherent incompleteness of information in RDF data sources and enrich query answers with implicit information using ontologies defining the underlying RDF data (Glimm & Ogbuji, 2013; Kostylev & Grau, 2015).

To infer a new statement, entailment regime defines a set of conditions (also called inference constraints) to represent a SPARQL query as *graph patterns* including a set of RDF triples with variables. Evaluating the graph patterns over the RDF data will return the answers satisfying the variables in the query. To implement SPARQL entailment regime, query rewriting techniques offer the means to achieve an inference by transforming/extending the query pattern and infer new statement that is not explicitly included in the RDF data (Jing, Jeong, & Baik, 2009).

However, current Semantic Web languages are based on classical logic (i.e., OWL that is based on description logics), which are deductive, monotonic and deterministic. The Semantic Web technologies, as opposed to the requirements of exploratory data analysis, are known to be inadequate for representing and reasoning with uncertainty and incompleteness (Jain, Gupta, & Bhardwaj, 2018; Predoiu & Stuckenschmidt, 2008; Redavid, Iannone, Payne, & Semeraro, 2008).

2.3.2 Reasoning with uncertainty

Uncertainty is an intrinsic characteristic of various types of applications that include data (and knowledge) processing, which significantly limits their capability. Especially in the medical context, uncertainties almost always outweigh what is known; patients and health care providers usually have little information when starting the diagnosis and treatments (Krumholz, 2014).

The term *uncertainty* can describe various forms of imperfect knowledge, including incompleteness, inconclusiveness, vagueness, ambiguity, or any other situation in which the Boolean truth values are unknown, unknowable, or inapplicable (Chen, Xiong, Yan, & Wang, 2018; Laskey et al., 2008). Uncertainties and lack of data manifest non-deterministic relationships (Zhu, Qu, Zhao, Chen, & Jalii, 2017). Bayesian and Fuzzy approaches, for a best guess estimation, have been commonly used to address the uncertainty within the (semantic) knowledge bases (Almeida, Kaymak, & Sousa, 2010).

Bayesian models leverage Probability theory to offer representation and reasoning systems for uncertain, incomplete knowledge. In Probability theory, truth-values are assigned to propositions identifying degrees of likelihood of a proposition that may be either true or false; truth-values range from zero (certain falsehood) to one (certain truth). *Bayes Rule* exploits truth values to calculate the probability of a proposition (event) based on prior knowledge.

Within the Semantic Web framework, Web Ontology Language (OWL), underpinned by Description Logics (DLs), is the main language for representing and reasoning over ontologies. However, the capabilities of DL, and inherently of OWL, will be challenged when it comes to the domains endowed with uncertainty. Probabilistic extensions to DLs and OWL allow for representing probabilistic ontologies and reasoning with them.

Probabilistic extensions of OWL (i.e., PR-OWL) (Ausín, López-de-Ipina, & Castanedo, 2014; R. Carvalho, Laskey, & Costa, 2013; R. N. Carvalho, Laskey, & Costa, 2017; P. da Costa, Laskey, & Laskey, 2008; P. C. Costa, 2005; Klinov & Parsia, 2008) are upper ontologies enabling the use of Bayesian theory for repressing and reasoning with uncertainty in the Semantic Web. PR-OWL aims to improve OWL capabilities for representing probabilistic ontologies based on Multi-Entity Bayesian Networks (MEBNs). This probabilistic extension to OWL can be considered as a bridge that connects deterministic ontologies defined in OWL with non-deterministic, probabilistic semantics of PR-OWL. However, in applications with large assertive databases, PR-OWL encounters some scalability issues due to the time complexity of OWL 2 DL reasoners to solve complex expressions (dos Santos, Carvalho, Ladeira, Weigang, & Mendes, 2015).

Extensions of OWL with fuzzy set theory (i.e., f-OWL) (Bobillo & Straccia, 2011; Liu, Huang, & Lin, 2013; Quach & Hoang, 2018; Stoilos, Stamou, Tzouvaras, Pan, & Horrocks, 2005) enable ontologies to capture, represent and reason with imprecise and vague information. The fuzzy OWL extensions are equipped with terminologies and grammars that are able to reason over fuzzy concepts, i.e., concept assertions, role assertions, concept inclusions, and role inclusions are associated with a degree of truth in [0,1], rather than a binary truth value (0 and 1) (Laskey et al., 2008).

Nikolaou et. al. (Nikolaou & Koubarakis, 2016) addressed the incompleteness issue in the Semantic Web frameworks which rely on the Open World Assumption (OWA). Nikolaou et. al. developed the semantics for an extension to RDF (RDF^i) to define new kind of literals (called *e-literals*) for each data type with the ability to represent property values that exist but are unknown or partially known using constraints. RDF^i extends the concept of an RDF graph by expressing partial information using the e-literals via a quantifier-free formula of a first-order *constraint language* $\mathcal{L}$. Nikolaou et. al. demonstrated the usefulness of the framework in geospatial Semantic Web applications. However, to support the constraints of $\mathcal{L}$ (e-literals), RDF^i requires the extension of FILTER operators in SPARQL queries, so that the e-literals can also be reflected in the condition part of the query and be evaluated over RDF^i databases. Hence, the management of e-literals makes the query manipulation complicated and impose some limitations as the e-literals are only allowed to appear in the object element of triples.

Agibetov et al. (Agibetov, Jiménez-Ruiz, & Solimando, 2015) proposed an evidence-based hypothesis testing method in the biomedical domain to cope with missing knowledge. They extract a causal chain from an ontology and represent it as Directed Acyclic Graph (DAG), which guides domain experts in conducting the experiment. This method only notifies physicians what knowledge is missing for the hypothesis to hold and does not generate any new knowledge.

The works studied here, and some similar works, are recent endeavors addressing the challenges with handling uncertainty and incompleteness in the Semantic Web framework.

Probabilistic approaches address the representation and reasoning with degrees of uncertainty about ambiguous pieces of information, and fuzzy formalisms allow for representing and processing degrees of truth about vague pieces of information (Laskey et al., 2008). Although these approaches are effective in the domains in which the data is joint with (quantitative) uncertainty and/or vagueness (or imprecision) (Lukasiewicz, 2017), neither of them can handle the uncertainty resulting from incompleteness. In addition, despite the effectiveness of Bayesian and Fuzzy methods to handle uncertainty, these methods need prior knowledge, including expert's input, statistical associations, or probability distributions, that might not always be available. These limitations challenge the efficiency of these approaches in coping with the rapid growth of data in different domains.

So far, the solutions for incompleteness have focused on two aspects: representation of missing data and finding the missing data. The latter approaches proposed extensions to knowledge representation formalisms providing the capabilities to capture and represent incomplete data (e.g., RDF[i]). One main motivation for representing the missing data is to help the analysis of data continue when it reaches missing data. The other type of approaches helps the user understand what pieces of data or knowledge are missing. Although, both approaches are helpful to manage knowledge incompleteness, they do not offer reasoning with incompleteness to derive the missing knowledge/data and *discover* the knowledge that is not captured and represented in any form of domain knowledge (i.e., ontology or rules). As such, we argue that there is still a lack of systematic, built-in support for ampliative, non-deductive Semantic Web reasoning that should be studied.

2.3.3 Query Rewriting within the Semantic Web

In the context of the Semantic Web, *query* implies the techniques and protocols that programmatically conduct information retrieval from the Web of Data, like data represented in RDF. SPARQL, the standard query language of the Semantic Web, is basically a graph pattern-matching query language that indicates how to construct the answer to a query (W3C, n.d.).

Rewriting a query regarding the ontological constructs allows the extraction of both explicit and implicit knowledge from the underlying data. In fact, Query Rewriting (QR) helps to ask what a knowledge base knows and what it assumes. Therefore, query rewriting is a technique to solve queries over an incomplete knowledgebase, which is supported with a set of terminological constructs (Grimm & Motik, 2005; Pérez-Urbina & Rodrıguez-Dıaz, 2012)..

In the query rewiring techniques, using DL-based constraints, a given query will be converted to a transformed query with regard to the relevant knowledge in the ontology (TBox), and then will be evaluated over the extensional knowledge about individual objects (ABox), as TBox-compliant statements about the ontology. For a given query Q and terminology $\mathcal{T}$, $Q_{\mathcal{T}}$ is the rewritten version of Q with regard to $\mathcal{T}$ (Figure 1), in which for every assertional knowledge $\mathcal{A}$, the solutions to Q over $\mathcal{T}$ and $\mathcal{A}$ can be obtained by evaluating $Q_{\mathcal{T}}$ over $\mathcal{A}$ (Pérez-Urbina, Motik, & Horrocks, 2010; Pérez-Urbina & Rodrıguez-Dıaz, 2012).

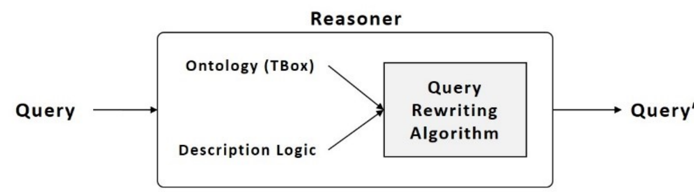

Figure 1- Query rewriting mechanism

The description logic profile underlying the query-rewriting algorithm identifies the inference power and complexity of rewritten queries. Based on the application and required level of expressivity, the query-rewriting algorithm can be loaded by different fragments of DL to meet the requirements of the application. Therefore, the flexibility of query rewriting in working with different profiles of DLs has made it a suitable approach for the cases where different level of ontology expressivity is required.

Query rewriting is appropriate for applications in which changes in the assertional data (ABox) occur more often than changes in the ontology (TBox). Additionally, reasoners implementing query rewriting can avoid large knowledge bases by extracting both explicit and implicit knowledge from the underlying data. In this regard, query rewriting potentially requires less memory and storage space (Pérez-Urbina & Rodrıguez-Dıaz, 2012).

As discussed earlier, within the SW framework, query rewriting techniques offer the means to implement SPARQL entailment regime to realize an inference by transforming/extending the query pattern and inferring a new statement that is not explicitly included in the RDF data (Jing et al., 2009). OWL 2 QL profile supports query rewriting mechanisms to explore data through domain knowledge. OWL 2 QL is underpinned by DL-Lite family of description logics. The Open Word Assumption (OWA) made in DLs makes query rewriting based on OWL 2 QL an appropriate tool to develop exploratory data analytics in the Semantic Web scenarios (Bienvenu, 2016; Grimm &

Motik, 2005). In the graph representation of data in RDF, query rewriting manipulates graph patterns in queries (i.e., SPARQL query) to explore the graph and discover hidden associations in the underlying data (Jing et al., 2009). Hence, query rewriting within the Semantic Web can be considered as a (description) logic-based query answering technique that leverages OWL constructs to explore large data represented in RDF graphs and derive new relationships among data.

Query Rewriting Approaches

Rewriting a query can be done by (i) *augmenting* the query by adding new terms and conditions retrieved from ontologies to the query, (ii) *trimming* or *relaxing* the query by removing the terms and pruning the conditions, or (iii) *substituting* the query terms with semantically related terms.

Logical relaxation approaches (Calì, Frosini, Poulovassilis, & Wood, 2014; Fokou, Jean, & Hadjali, 2014; Huang, Liu, & Zhou, 2012; Hurtado, Poulovassilis, & Wood, 2008) relax a query by manipulating the conditions of the query — e.g., by replacing constants with variables or by using the class and property hierarchies based on RDFS entailment and RDFS ontologies. Query rewriting approaches (Dolog, Stuckenschmidt, Wache, & Diederich, 2009) rewrite the query based on some predefined rules. Each rewriting rule consists of a matching pattern, a replacement pattern and a set of conditions that restrict the applicability of the rule. A query rewriting method will replace the matching pattern in the original query with the replacement pattern in the new query, whenever the conditions are met. Moreover, statistical language models (Elbassuoni, Ramanath, & Weikum, 2011), matching functions (Hogan, Mellotte, Powell, & Stampouli, 2012), and *failure causes* detection models (Fokou, Jean, Hadjali, & Baron, 2016) are other RDF flexible querying techniques.

In addition to the computation time of the relaxation process, query relaxation approaches reduce the query constrains with the aim of making it possible to retrieve answers with varying degrees of exactness, and do not perform any knowledge discovery or data analytics techniques on the RDF data. These approaches are beneficial only in the cases where the query contains avoidable constraints and conditions (over-constrained), e.g., users usually provide very broad queries or are often not able to correctly formulate queries. Hence, query relaxation is not efficient for the simple queries that contains no or few constraints, and still the KB cannot provide any answer for them.

Augmented queries and trimmed queries, whether they are *conjunctive* (AND) or *disjunctive* (OR), return a super-set or sub-set of the original query. Whereas, substitution explores the ontology, replaces the concept(s) in the query with semantically relevant entities and may return a new result set, which may partially overlap the original result set (Mangold, 2007). This exploratory technique is beneficial to the cases where the augmentation and trimming yield no answers.

When representing the rewritten queries, the approaches above may differ from each other (Venetis, Stoilos, & Stamou, 2014). *Union of Conjunctive Queries (UCQs)* is the most common approach of representing a so-called perfect rewriting of a query. A union of conjunctive query Q is a set of conjunctive queries of the same arity and having the same query predicate. However, UCQ is not a golden key. Depending on the level of ontology expressivity that would be considered in the query-rewriting algorithm, it might turn out a query that is too big or too complex, which would compromise the feasibility of its evaluation. This means a UCQ might contain hundreds or thousands of queries that its evaluation on the ABox is exponential w.r.t the size of the original query. One approach to reduce the size of union of conjunctive queries is query optimization by considering ABox, while in applications with large ABoxes this approach has its own drawbacks.

Besides UCQs, *Datalog Queries (DQs)* is a set of rules, such that the head predicate of each rule is not a predicate used in the ontology. In comparison to UCQs, these types of query are harder to evaluate, but it is a solution to the cases that further ontology expressivity is required (Pérez-Urbina & Rodrıguez-Dıaz, 2012; Rosati, 2012).

2.4 Knowledge graph

Knowledge Graph (KG) is a graph-based knowledge representation formalism in the big data era, organizing massive volumes of data from multiple sources in multiple topical domains. Knowledge graphs can be considered as a set of big semantic networks connecting the real-world entities and their relationships in the form of predicate-argument structures, e.g., *subject-predicate-object*. Currently, large knowledge graphs in different areas, such as DBPedia, Yago and Google knowledge graph, supply a large amount of data and knowledge for the use of both the research community and the commercial sector (Gunaratna, 2017; Paulheim, 2017; Sadeghi, Lange, Vidal, & Auer, 2017).

Knowledge graphs capture the relationships between the concepts and connect fragmented pieces of knowledge together. Knowledge graphs encode semantic associations through capturing different types of entities (i.e., nodes), their properties (i.e., arcs) and the chain of relationships that connect those entities, providing the opportunity to reveal interesting and unknown connections between different entities (Anyanwu & Sheth, 2003; Bianchi, Palmonari, Cremaschi, & Fersini, 2017; Yang, Huang, Han, Hua, & Tang, 2017). In particular, the semantics encoded in the knowledge graphs enrich graph structure (and graph database) into something greater than a graph-based knowledge representation reinforcing wider range of services (e.g., querying, analytics, reasoning, semantic search, etc.) than a plain graph database offers (Yang et al., 2017).

Knowledge graphs are distinguishable from more traditional knowledge representation formalisms, because they simultaneously provide *normalization* (i.e., information is accessible through smallest units, e.g., entities), *connectivity* (i.e., relationships between the units carry the knowledge) and *context* (i.e., connections are annotated with contextual information to include meta-data). Hence, knowledge graphs can be considered a helpful platform for the task of information retrieval, knowledge discovery and question answering (Krotzsch, 2017; Y. Zhang, Dai, Kozareva, Smola, & Song, 2017).

The aforementioned characteristics may associate knowledge graphs with knowledge-bases and ontologies. Although they are often used interchangeably, they are indeed different concepts. Knowledge graphs are different from ontologies with regards to two aspects: (i) quantity (size); knowledge graphs can be considered as large ontologies that contain not only classes and properties, but also include instances, and (ii) extended requirements; like a built-in reasoner allowing the extraction of new knowledge. Furthermore, knowledge graphs are superior to knowledge-based systems that consist of a knowledge base and a reasoning engine. A knowledge graph extends conventional knowledge-based systems with collection, extraction and integration of information from additional, external sources. Ehrlinger et. al. emphasize the reasoning capabilities of knowledge graphs and describe it as a knowledge-based system that employs information integration: "*A knowledge graph acquires and integrates information into an ontology and applies a reasoner*." (Ehrlinger & Wolfram, 2016; Sadeghi et al., 2017)

Furthermore, the Semantic Web technologies offer efficient knowledge representation, integration and reasoning formalisms by organizing numerous heterogenous data sources into knowledge graphs using Resource Description Framework (RDF), RDF Schema and Web Ontology Language (OWL) ontologies (L. Shi et al., 2017). Considering the layers of the Semantic Web framework, logic-based formalisms offer features that a knowledge graph may utilize:

a. A semantic representation of knowledge with various levels of granularity, which carries syntax of data while captures its semantics: RDF stores data in the triple format and constructs a flexible and extensible knowledge graph that facilitates the incorporation of newly discovered facts. RDFS and OWL ontologies represent associations between concepts, model distinct types of data, while seamlessly capture its semantics in an accurate, rich and unambiguous way;
b. Built-in support for deduction-based reasoning: ontology-based reasoning, rule-based reasoning and entailment regime of SPARQL 1.1 conform with the Open Word Assumption (OWA) and facilitate the generation of new knowledge by analyzing the semantically annotated data;
c. Linked Open Data (LOD) integrates external data sources to further extend the data coverage and to synthesize a unified large global data source,
d. Query languages are utilized to navigate, manipulate and retrieve the data stored in heterogenous, distributed graph structures (i.e., triple repositories).

Therefore, the Semantic Web technologies provide a modern and efficient platform to implement knowledge graphs that can crawl the Web, collect and process (i.e., analyse) the information, and deliver interesting and unknown insights (Ehrlinger & Wolfram, 2016).

2.4.1 Knowledge graph analytics

The analysis of large amount of (heterogeneous) data combined with its semantics (i.e., relationships between data items) in the form of a graph provides new opportunities to semantic data analytics, which traditional analytical approaches may not be able to offer. Additionally, the analysis of connected knowledge sources (in the form of knowledge graphs) offers a potential to explore independent disciplines (i.e., health care, public health, economics, political science, and etc.) and identify relationships that are impossible to be discovered within a single discipline alone (Roski et al., 2014). However, large knowledge graphs, which include numerous entities along with their relationships, pose great challenges to the traditional logic-based reasoning systems. Hence, there is a need for innovative analytical approaches that can exploit graph representation of knowledge and offer predictive and real-time analysis (Chen et al., 2018; Roski et al., 2014; Wei, Luo, & Xie, 2016).

Velampalli et. al. (Velampalli & Jonnalagedda, 2017) proposed a framework to extract common skill-sets from resumes. They introduced a MapReduce algorithm to model the skill-sets from resumes into a graph structure (conceptual graph). Afterwards, they identify common skill-sets from the resulting graph using SUBDUE (Holder, Cook, & Djoko, 1994), a popular Graph Based Data Mining (GDM) method. Their experimental results showed the efficiency of graph mining algorithms in the skill-set analytics. The framework introduced by Velampalli et. al. and other GDM techniques, such as AGM (Inokuchi, Washio, & Motoda, 2000), Gaston (Nijssen & Kok, 2004), etc., look for common and frequent substructures, similarities, and anomalies in conceptual graphs. Their exploratory closure of these approaches is limited to the patterns that are frequently occurred in the graph. Hence, they are not able to discover the patterns that exist but are not frequent enough to be detected by the graph mining algorithms.

Zhao et. al. (Zhao, Munne, Kertkeidkachorn, & Ichise, 2017) analyze knowledge graphs to discover missing RDF triples. To this end, they apply string-based similarity measures and Recurrent Neural Network (RNN) to identify similar entities in two knowledge graphs (e.g., DBpedia and YAGO). They perform a graph-based ontology integration method for mapping the similar entities. In their approach, they successfully discover missing and incorrect RDF triples using the Semantic Web technologies and Natural Language Processing (NLP) techniques. In their approach, they always require (at least) two knowledge graphs (i.e., ontologies) to match the existing triples, identify the missing ones, and discover the correct RDF triples. This limitation restricts the applicability of the approach to the cases that only one incomplete knowledge graph is available.

Shi et. al. (L. Shi et al., 2017) introduced an automatic healthcare knowledge retrieval system from textual medical knowledge. They build a structured ontological model of health data from Electronic Health Record (EHR) systems, organize medical text into conceptual graphs, and provide semantic mapping between medical text and medical knowledge. Their model is stored in relational databases. Afterwards, the framework explores complex semantics between the entities in the graph to automatically retrieve knowledge. Shi et. al. performs contextual inference to prune the knowledge graph and avoid the meaningless results, but this inference does not extract the knowledge hidden in the complex relationships.

Yang et. al. (Yang et al., 2017) proposed an approach utilizing knowledge graphs to analyze the neglected influencing factors of statin-induced myopathy in a case of the coronary heart disease. In their approach, they generate semantic SPARQL queries based on the patient's history, symptoms, blood test results, etc. to search Linked Life Data[1], the knowledge graph. The resulting query will return a set of relevant documents helping to find relationships between the disease, their symptoms and side-effects. Yang et. al. showed the effectiveness of knowledge graphs in answering the medical cases that involve various concepts and parameters. However, the process of generating an optimal SPARQL query requires some background knowledge about the SPARQL query language and the structure of the knowledge graph—e.g., the relevant UMLS concept ID or URI of the concepts in the Linked Life Data. This requirement limits the applicability of their approach in real medical settings.

The works above show the utility of knowledge graphs in real-world applications. Knowledge graphs construct an integrated platform for the management of a massive volume of data and knowledge, while they provide the opportunity for semantic data analytics and knowledge discovery. These studies demonstrate knowledge graphs can capture data, domain-specific knowledge and the semantics of data. In applications (e.g., healthcare) in which the association between data items (e.g., disease-treatment interaction) is what is of most importance, the analysis of knowledge graphs provides a new way of thinking, training, and method to find interesting characteristics and actionable

[1] http://linkedlifedata.com/

insight from underlying data (Gunaratna, 2017; Krumholz, 2014; Miller, Ramaswamy, Kochut, & Fard, 2015).

2.5 Summary

This chapter introduced semantic analytics as an exploratory approach that explores semantically annotated data to derive actionable insights from (large) data facilitating decision making and problem solving. The related work showed the exploitation of data semantics can significantly enhanced data interpretation, especially by non-data experts. However, in practice, there still exists a lack of mechanisms that implement reasoning approaches over large semantic data and discover complex, nuanced insights needed for making better informed decisions (S. R. Abidi, Cox, Abusharekh, Hashemian, & Abidi, 2016; S. S. R. Abidi, 2001; S. S. R. Abidi & Hussain, 2007).

Along with the requirements of semantic analytics with respect to data representation and analysis, the capabilities of the Semantic Web technologies were studied. It has been seen that the Semantic Web languages provide expressive knowledge representation formalisms that seamlessly capture and represent syntax and semantic of data. They also allow the generation of new knowledge by analyzing the underlying semantics of the available knowledge through a variety of reasoning mechanisms. It is shown that the graph representation of large data using RDF graphs and ontology languages (RDFS and OWL), encodes the semantic associations between entities and offers new ways of analysis to find actionable insight from the underlying data.

However, the expressivity of the Semantic Web ontology languages is limited—especially when it comes to model complex domains with large, diverse data. This chapter studied several studies that successfully introduced extensions to the Semantic Web languages providing the required expressivity to represent complicated associations between data items as well as the uncertainty associated with the data. The novel approaches capable of working with the extensions, analyzing the semantic data and reasoning with uncertainty were studied. But there is still a lack of exploratory data analytics approaches that leverage semantic data and infer new knowledge from large data-sets when dealing with incompleteness.

It is discussed that query rewriting techniques offer the means to derive insights from data, which are not already stated in the data or domain knowledge. Nevertheless, a query rewriting engine requires a set of conditions (aka. inference constraints) conducting the rewriting. In the next chapter, we introduce plausible reasoning as a non-deductive, exploratory data-driven reasoning approach that can provide the foundations of developing an ampliative query rewriting engine.

Chapter 3: Plausible Reasoning

In this chapter we study plausible reasoning as a weak form of inference that leverages the semantics of relevant concepts and allows for dealing with incomplete data during decision making. We also investigate the implementation of plausible reasoning through a set of plausible patterns (e.g., generalization, specialization, interpolation, etc.), that are applied to a variety of semantic relationships (e.g., conceptual hierarchy, partial order, etc.). But before studying the components of plausible reasoning, we need to understand what the *plausibility* is and what are the circumstances that plausible reasoning can contribute to.

3.1 The notion of plausibility

Plausibility implies the situations featuring some degree of imprecision that is not quantified. A hypothesis is plausible when the evidence supporting the hypothesis is stronger than the arguments against it. So, for the time being, the hypothesis is (plausibly) acceptable (Cellucci, 2013a). However, an accepted plausible claim can turn out to be false, and a implausible claim can end up being true (Bunnin & Yu, 2004).

From Alexander's perspective (*In Top. 19.22–27*), the Greek philosopher, plausibility is different from being true, not by being false, but by the principles on which the judgement is based upon—i.e., even though a plausible opinion can be in fact true. The truth of plausibility is not only dependent on things, but the judgment involves the listeners and their assumptions about things as well (Vega Renon, 1998). Aristotle (384–322 BC) confirms "the man who makes a good guess at truth is likely to make a good guess at what is plausible." (*Rh. 1355a15–18*). He believes in a set of plausible opinions/arguments "it is acceptable that any of these discoveries may be entirely wrong, but rather that they should be right in at least some respect or even in most respects” (*EN1098b27–29*).

Renon (Vega Renon, 1998) leverages Aristotle’s conception on plausibility to characterize plausibility in two basic features: plausible opinions are (i) pragmatic in nature, and (ii) graded. The first feature emphasizes that *plausibility* is not a semantic property, but a practical (*pragmatic*) relationship that is recognized effective by a person, group or community. And because of its pragmatic nature, a plausible opinion is rational (*dialectical*) as well.

In the absence of deterministic, exhaustive knowledge, plausibility is akin to natural experiments, when the investigator leverages the available data to make assumptions, examine possible outcomes and present alternatives (Habicht et al., 1999). Plausibility assumes there might exist more information that we are not aware of. That assumption conforms to the Open Word Assumption (OWA) in which a non-existing fact is not assumed as false, but as unknown.

Contrary to probability, there are no numbers, e.g., degrees or level of certainty that can quantify the lack of precision in plausibility. However, plausible situations can be associated with frequency of occurrence (i.e., indicating the cases that something is true more often than being false) or weight of evidence (i.e., when the evidence for something outweighs the evidence against it) (Billington, 2017).

Plausible opinions can be graded based on their weight and authority; a more plausible opinion has greater degree of real acceptance. However, although plausible opinions can be compared amongst each other (i.e., regarding their degree of plausibility), they would exclude the possibility of mutual or internal conflicts— e.g., opinion α, accepted by a community, is more plausible than β, the opinion of the expert; however, β is still relatively plausible even though it might contradict opinion α (Vega Renon, 1998).

3.2 Plausible Reasoning

Plausible reasoning relates to the human's problem-solving process, which analyzes the semantics of the available data to discover unknown associations inherent within the data and establish meaningful relationships to solve complex problems. Plausible reasoning provides a non-deductive exploratory approach to solve problems when there is enough evidence to justify the *plausibility* of the solution(s) through semantic analysis of data.

As an exploratory analytical method, plausible reasoning has the capability to discover new information/knowledge. Conventional (deductive) reasoning approaches reason with a *complete* set of true statements to derive another true statement that was already contained within the premises. But plausible reasoning explores a partial set of true statements (i.e., incomplete data) to derive a *plausibly true* inference, which is the best-effort answer in light of what is known so far. This new inference is often reasonable when it is used under the right conditions (Cellucci, 2013b; Jaynes, 2003; Mohammadhassanzadeh, Van Woensel, et al., 2017; Walton, Tindale, & Gordon, 2014).

3.2.1 Characteristics of plausible reasoning

With regard to the characteristics of human plausible reasoning, Tindale (2010) and Walton (2013) identified eleven characteristics for plausible reasoning:

1. Plausible reasoning proceeds from premises that are more plausible to a conclusion that was less plausible before the plausible argument.

2. Something is found plausible when the audience has examples in their own minds.
3. Plausible reasoning is based on common knowledge.
4. Plausible reasoning is defeasible (non-monotonic).
5. Plausible reasoning is based on the way things generally go in familiar situations.
6. Plausible reasoning can be used to fill in the implicit premises in incomplete arguments.
7. Plausible reasoning is commonly based on appearances from perception.
8. Stability is an important characteristic of plausible reasoning.
9. Plausible reasoning can be tested, and by this means, confirmed or refuted.
10. Probing into plausible reasoning in a dialogue is a way of testing it.
11. Plausible reasoning is graded, but it is different from the standard probability values and Bayesian rules used in Pascalian probability.

The third characteristic puts an obligation on a plausible inference that the sequence of argumentations leading from acceptable evidences to ultimate inference should be (nearly) known by everyone. It means that the chain of inference used to derive a plausible solution should be justifiable and admissible—i.e., plausible reasoning conducts an inference from some generally accepted evidence (characteristic 2) to new features that are (somehow) familiar (characteristic 5).

Plausible reasoning is particularly useful in the situations in which the data/knowledge is incomplete—i.e. the reasoner is expected to derive a conclusion while not having a complete set of knowledge. In this regard, plausible reasoning is ampliative and non-demonstrative (characteristic 6), as it extends the knowledge by learning and discovering new pieces of knowledge, based on what is known so far.

The eighth characteristic implies the consistency among a set of plausible hypotheses. Plausible reasoning is stable as an argument can be strengthened or weakened by another. For example, a physician concludes that a man has fever not only from one symptom (such as rapid pulse or high temperature) but also from other existing symptoms (such as soreness of touch or thirst)—i.e., each symptom is consistent with what the other symptom(s) suggest. Based on the characteristics above, plausible reasoning can be recognized as an inference method that is:

- **Non-demonstrative;** non-demonstrative reasoning methods depend on knowledge discovery, making hypotheses, and learning new concepts. Typical examples include medical diagnosis, economical statistical evidence, or findings of a scientific research. On the other hand, demonstrative reasoning (e.g., mathematical proof) is sound, deterministic, beyond controversy, and final, but it is incapable of exploring new knowledge (Pólya, 1954).
- **Ampliative;** non-ampliative reasoning, like deduction, explicates and instantiates what was already expressed in the captured domain-specific knowledge (e.g., via deductive rules). While, ampliative reasoning generates inferences that go beyond what is contained (known) in the captured knowledge and explores what it assumes (Blachowicz, 1989; Ippoliti, 2008).
- **Non-monotonic;** a plausible conclusion can be retracted in the light of further information. Non-monotonic logics have been devised to overcome the limitations of classical logics to capture and represent *defeasible inference.* Despite classical (monotonic) logics, in which the inferences are deductively valid, non-monotonic reasoning draws an inference from a set of facts, knowing that new facts may

challenge the previous conclusions (Antonelli, 2008; Billington, Estivill-Castro, Hexel, & Rock, 2006; Nute, 2001).

- **Subjective;** a plausible argument is an expression of beliefs, opinions, personal preferences, values, feelings, and judgments (Jøsang, 1997).

3.2.2 Plausible reasoning is different than probabilistic reasoning

Mutually exclusive solutions are a set of possible responses to a question or a situation in which the correctness of one solution decreases the validity of the other solution(s). In other words, in a set of feasible mutually exclusive solutions that exhaust the possibilities (i.e., the sum of the probabilities is equal to 1), it is not possible that two or more solutions will be simultaneously acceptable. Although in some scenarios, the outcome of inference is mutually exclusive by nature, there are some domains (such as health, law) that multiple (sometimes contradicting) solutions are feasible at the same time.

In probabilistic reasoning, the probability of the correctness of a solution is equal to 1 minus the probability of correctness of other mutual exclusive solution(s)—i.e., the probability of a statement *not-A* is calculated as *1-pr(A)*. While, plausibility does not involve the statistical likelihood of the possibilities and solutions are not necessarily mutually exclusive—i.e., a new fact alters the evidence, and detracts or strengthen the plausibility of a solution (Walton et al., 2014).

3.2.3 Plausible reasoning is different than fuzzy logic

There are two types of uncertainty that have different impacts on reasoning and knowledge representation. The first type of uncertainty (ς_1) occurs whenever the knowledge is inherently entangled with some sort of ambiguity or contradiction. The other type of uncertainty (ς_2) is a result of missing knowledge (i.e., incompleteness) (Dompere, 2012).

There are different approaches to deal with these two types of uncertainty. Plausible reasoning provides methods to overcome the uncertainty resulting from incompleteness and lack of knowledge (ς_2). These approaches, like induction, analogy, and abduction acquire new knowledge through learning and discovery. On the other hand, statistical methods like probabilistic reasoning, fuzzy logic, Bayesian belief network, Dempster-Shafer models try to solve the first kind of uncertainty (ς_1). Statistical methods provide insight into the likelihood of certain items and represent beliefs that are not certain (Han, Klein, & Arora, 2011).

3.3 Semantic associations applicable to plausible reasoning

A variety of semantic relationships plays a significant role in plausible reasoning. In particular, plausible reasoning relies on fine-grain knowledge representing how different concepts are semantically associated (Derrac & Schockaert, 2015). Three types of semantic relationships (NISO Standard (ANSI), 2010) may be applicable to plausible reasoning: hierarchical relationships, equality relationships, and associative relationships.

3.3.1 Hierarchical relationships

Hierarchical relationships (Table 1) imply class inclusion and are based on degrees or levels of super-ordination and sub-ordination. The superordinate term represents a class or a whole, and subordinate terms refer to its members or parts. Hierarchical relationships include three different associations: generic relationships (is-a relationship, such as class-member relationship), instance relationships (is-a relationship, such as class-instance relationships), and whole-part relationships (such as part-of relationship).

Table 1- Hierarchical semantic relationships

Semantic relationship	Symbol
Element of/Part of	$\in$
Contains (as instance)	$\ni$
Subset of	$\subset$
Superset of	$\supset$

3.3.2 Equality relationships

Equality relationships (Table 2) express the associations between two or more variants of a same concept, which are equivalent or nearly-equivalent. An equivalence relationship may belong to one of the basic types of synonyms, lexical variants, near-synonyms, generic posting, or cross reference to elements of compound terms.

Table 2- Equality semantic relationships

Semantic relationship	Symbol
Equal to	$=$
Almost equal to	$\approx$
Not equal to	$\neq$
Identical to	$\equiv$
Similar to	$\sim$

3.3.3 Associative relationships

Associative relationships (Table 3) represent associations between related concepts, which are neither hierarchical nor equivalence, but still semantically or conceptually related. Associative relationships cover several types of relationships: *cause/effect* (such as infection/hearth failure), *action/product* (such as classification/order), etc. Some varieties of associative relationships can be interpreted using (partial) order theory, in which concepts are being ordered in an interpretable direction regarding a measurable property such as size, location, time, etc. Hence, any (binary) associative relationship can be considered as a partial order association if it is:

- Reflective ($a \leq a$),
- Asymmetric ($a \leq b, b \leq a\ then\ a = b$), and
- Transitive ($a \leq b, b \leq c\ then\ a \leq c$).

Table 3- Ordered semantic relationships

Semantic relationship	**Symbol**
Less than	$<$
Greater than	$>$
Precedes	$\prec$
Succeeds	$\succ$

3.3.4 Representation of plausible semantics

Plausible reasoning depends on the comprehensiveness of the knowledge representation formalism that incorporates diverse semantic associations connecting various entities. In this regard, fundamental requirements of an efficient representation of the semantic associations applicable to plausible reasoning include (Anshakov & Gergely, 2010; Cohen & Conway, 2007; Collins & Michalski, 1989; Davis, 1990; Halford, Wilson, & Phillips, 2010; Kuipers, 1979; Panton, Matuszek, & Lenat, 2006):

- Delivering required data/knowledge expressivity
- Representing conceptual hierarchical relationships
- Representing ordered relationship
- Ability to handle complex (both syntactically and semantically) structures

- Capability to express context-aware knowledge
- Representing intermediate states of reasoning process (e.g., arbitrary relationships)
- Capability to represent ‘don’t know’

Nevertheless, not every knowledge representation formalism automatically supports all the requirements above. This limitation imposes some challenges for capturing and representing the plausible semantics and consequently restricts the effective implementation of plausible reasoning. The challenges associated with the representation of the associations applicable to plausible reasoning is elaborated in this chapter after introducing the plausible patterns. Our approach to address the challenges is discussed in the next chapter.

3.4 Plausible patterns

Plausible reasoning leverages semantic associations to perform inference through a set of frequently recurring patterns and suggest a plausible statement, which could be further tested deductively (Cellucci, 2013b; Collins & Michalski, 1989). These patterns explore the semantic associations between the entities and do not occur in classic forms of logic (Mohammadhassanzadeh, Van Woensel, et al., 2017; Virvou & Kabassi, 2004).

To identify the recuring inference patterns of plausible reasonign, Collins and Michalski (Collins & Michalski, 1989) collected considerable number of people’s answers to common questions. They studied how people connect together different pieces of knowledge to draw an answer that they didn’t know beforehand. They realized there are different inference patterns that are used to answer a question, while same patterns appear in different reasoning processes. The analysis of the answers resulted in the identification of a taxonomy for plausible patterns: *generalizaiton*, *specialization*, *similarity* and *dissimilairy*. All these patterns leverage hierarchical realtionhsips to conduct an inference.

Later, Cellucci (Cellucci, 2013b) introduced these patterns along with *metaphor*, *metonymy*, *definition as analysis* and *diagrams* as the rules of discovery. He considers the patterns as non-deductive rules, which develop solve poblems via generating plausible hypotheses using the already available data. In his theory, the ampliativity of the rules is paramount. Cellucci considers a *'heuristic power'* for the rules, which discover a conclusion that is not contained in the premises. In a different study (Derrac & Schockaert, 2014, 2015), Derrac and Schockaert investigated *interpretable directions* between obejcts in conceptual spaces to infer plausible conclusions. They formulated *interpolation* and *a fortiori* as two reasoning patterns that conduct commonsense reasoning to fill the gaps in knowledge bases.

Along with the characteristics of plausible reasoning, the studies above emphasize on the *ampliativity* of the plausible patterns—i.e., the inference patters can *discover* conclusions that were not included in the premises (non-demonstrative). However, they confirm that the plausibility of findings is influenced by the available data and is approved for the time being (non-monotonic).

In this section, we introduce eight plausible patterns. Later, we discuss the feasibility of implementing the patterns with respect to the semantic associations and formalizing the rationale behind the patterns. However, we realize that the plausible patterns are not limited to the patterns listed here. More patterns may exist in the literature (i.e., philosophy, law, etc.) that were not identified here.

3.4.1 Generalization

Generalization is passing from a given set of objects S, to a larger set S', that contains the given set ($S \in S'$). In hierarchical structures, generalization involves moving from one node (concept) to its parents. For example, *daffodil* is a type of flower and *England* is in Europe. A statement like *flower_type(England, daffodil)* can be generalized to *flower_type(Europe, daffodil)*, since *Europe* includes *England*, and consequently includes whatever grows in it (Cellucci, 2013b; Collins & Michalski, 1989; Heit, 2000). The implication is if $b \subseteq a$ and A is true about b, then A is true about a as well:

$$Generalizaiton: \frac{b \subseteq a \qquad A(b)}{A(a)}.$$

3.4.2 Specialization

Specialization is passing from a given set of objects to a smaller set that is contained in the initial set. In hierarchical structures, specialization implies moving from one node (concept) to its children. Regarding the prior example, the statement can be transformed to *flower_type(London, daffodil)*, since *London* is a part of *England* and *daffodil* (plausibly) grows in *London* as well (Cellucci, 2013b; Collins & Michalski, 1989; Heit, 2000). The implication is if $a \subseteq b$ and A is true about b, then A is true about a as well:

$$Specialization: \frac{a \subseteq b \qquad A(b)}{A(a)}.$$

3.4.3 Interpolation

Interpolation is a mapping from observation space X to conclusion space Y, where $x_i \in X$ is not explicitly mapped to any y $\in Y$, while there are direct mappings entities from space X to space Y which are relevant to x and y. Interpolation is based on the qualitative notion of conceptual betweenness. For example, we know that *undergraduate* and *PhD* students are both exempt from paying council tax in the UK. Hence, it is *plausibly inferable* that *master's* students are also exempt from paying this tax, knowing that master's students are conceptually between undergraduate students and PhD students (Derrac & Schockaert, 2015). The implication is if $a \leq b \ \leq c$ and A is true about a and c, then A is true about b as well:

$$Interpolation: \frac{a \leq b \ \leq c \quad A(a) \quad A(c)}{A(b)}.$$

3.4.4 A Fortiori

A fortiori argument is an inference from a proposition with high degree of confidence to a less confident proposition that is not explicitly specified but it is implicit in the first one. Based on the a fortiori argument, truth of a statement implies another statement, which is included in the first one and is plausible, common and familiar (Derrac & Schockaert, 2015; Zurek, 2012). A fortiori reasoning exploits the relationships between concepts, definitions and actions that have been expressed by means of a partial order. Depending on the orientation, a fortiori reasoning can be conducted in two variants: a *maiore ad minus* (from more to less) and *a minori ad maius* (from less to more) (Hallaq, 2009; Schockaert & Derrac, 2015; Sion, 2009).

More to less a fortiori indicates an inference from greater to smaller, general to particular, whole to part, and stronger to weaker. For example, if a door is big enough for a person with two meters high (greater case) to pass, then a shorter person (smaller case) can pass through as well. The implication is if b occurs after a and A is true about b, then A is true about a as well:

$$A\,fortiori - more\ to\ less\colon \frac{a \le b \quad A(b)}{A(a)}.$$

To the contrary, *less to more a fortiori* denotes making an argument in the reverse direction. For example, if a person knows buying beer is illegal under the age of 18, then it is plausibly inferable that buying whiskey is illegal since it includes more alcohol than beer. The implication is if b occurs after a and A is true about a, then A is true about b as well.

$$A\,fortiori - less\ to\ more\colon \frac{a \le b \quad A(a)}{A(b)}.$$

3.4.5 Similarity/Dissimilarity

Similarity or dissimilarity is moving between any two comparable nodes in the hierarchy. For example, *rose* and *daffodil* grow in temperate climates. So, they are conceived similar. But *bougainvillea* is different as it is a subtropical flower. In this regard, *flower_type(England, rose)* is a plausible statement transform, while *flower_type(England, bougainvillea)* cannot be a valid inference (Collins & Michalski, 1989).

In addition to the hierarchical equality, similarity and dissimilarity can be performed by considering equivalent concepts w.r.t. an ordered semantic. *Ordered-equivalence* semantics compare two concepts that overlap each other. For example, relationships like same age, same height, same severity, etc. imply a similarity between two concepts w.r.t. an ordered property. Likewise, *ordered-based dissimilarity* exploits the associations that express no overlap. An o*rdered-based dissimilarity* relationship can be considered as an *a fortiori*, in which we don't know the direction of the ordered relationship between the concepts. In similarity, the implication is if b and a are considered equal w.r.t. some properties (i.e., conceptual or measurable), and A is true about a, then A is true about b as well:

$$Similairty\colon \frac{a \sim b \quad A(a)}{A(b)}.$$

3.4.6 Metaphor

Metaphor assumes that entities of one domain (the target domain) are equivalent to the entities from another domain (the source domain), while there is no reasonable connection. Metaphor is considered as an ampliative pattern because it allows to find new hypothesis. A hypothesis driven by metaphor suggests if entities from the source domain have a certain property, then the entities from the target domain have that property as well. For example, if the target domain includes heart and the source domain consists of glass, then "Heart is made of glass" is a metaphor, implying people could be *emotionally* hurt and sad (Cellucci, 2013b).

Metaphor compares two not-related things and tries to find some pre-existing similarities, which may not necessarily mean the same in the original domains. For example, in the case of "Heart is made of glass", heart and glass are assumed similar in the sense of being fragile: while glass is breakable into pieces, heart cannot be physical hurt—i.e., heart is emotionally vulnerable (Cellucci, 2013b). The implication is if T and S are the target and source domains respectively, anything that is true about the elements of S is applicable to the elements of T as well:

$$Metaphor\colon \frac{T \longmapsto S \quad a \in T \quad x \in S \quad A(x)}{A(a)}.$$

3.4.7 Metonymy

Metonymy indicates using one thing instead of something else that is commonly associated with it. Metonymy is used in mathematical symbolisms where a letter (or a combination of letters) represents a triangle, or in a conversation when ‘Wall street’ denotes the New York Stock Exchange and its role in the international financial system (Cellucci, 2013b). Metonymy is different than metaphor. Metaphor is based on the similarity among the properties of two entities from two domains and suggest a new similarity. But metonymy draws an association between two concepts form a same domain—i.e., the association is stronger as more the two items are perceived together. The implication is if a stands for b, anything that true about a is applicable to b as well:

$$Metaphor\colon \frac{a \Rightarrow b \quad A(a)}{A(b)}.$$

3.4.8 Other plausible patterns

In addition to the patterns above, there are more plausible patterns (e.g., *definition as analysis*, *diagrams*, etc.) defined and utilized in different domains of scholastic studies. Although they are effective approaches to non-deductive reasoning, they could be treated as a subtype of one of the patterns above. For example, discovery via diagram, which resembles metonymy, allows a figure to represent another concept or entity. Or ‘definition as analysis’ attempts to find a common formal property and generalize the other properties. In this aspect, we found ‘definition by analysis’ a more specific version of similarity.

In addition, there exist some other patterns (e.g., *apagoge*) that are although considered non-deductive, there is a fundamental dispute among philosophers and mathematicians on their power of discovery (aka. ampliativity). For example, Peirce (Peirce & Ketner, 1992) believed abduction is what Aristotle meant by *apagoge*. While he once believed abduction is ampliative, later he admitted the idea of abduction being ampliative is weak. He argued an ampliative inference discovers something not implied in the premises, which is not true about abduction (Cellucci, 2013b).

3.4.9 Plausible patterns in scientific discoveries and reasoning

In addition to the common-sense examples that were provided to describe the rationality of the plausible patterns, they have also been extensively used in scientific discoveries in different domains all along the history of science.

Newton utilized generalization to infer that all the bodies have the power of gravity. His implication was that, there is a power of gravity pertaining to all planets (his earlier discovery) and since the planets are bodies, then "there is a power of gravity pertaining to all bodies."

Gauss, at the age of nine, leveraged specialization to solve the problem asking the sum of all numbers in the set $\{1,2,3, \dots ,100\}$. He solved the problem by considering the general problem, 'finding the sum of all numbers in the set $\{1,2,3, \dots, \mathrm{n}\}$.' Having that general problem solved ($sum = n(n + 1)/2$), by specialization, he inferred that the sum of numbers from 1 to 100 is 5,050 (Cellucci, 2013b).

Metaphor was used by Newton to discover a solution to the problem "Given any relationship whatever of fluent quantities, to find the relationship of their fluxions." He considered the growing (fluent) quantities as physical objects that are moving (i.e., mathematical quantities construct the target domain and the source domain is comprised of physical quantities). Using a bottom-up analytic approach to mathematics, Newton found that if two fluents x, y are in the relationship $y = x^2$, the ratio of their fluxions is equal to $2x$ (Cellucci, 2013b).

Metonymy has been widely used in geometric discoveries by Pythagoras to reason on concrete objects and then transfer his findings to abstract objects. For example, in the discovery of 'the three angles of a triangle are equal to two right angles', metonymy permits to reason over a drawn figure (i.e., the concrete object) and transfer the established properties to a triangle (i.e., the abstract object). Although it may seem that metonymy is not capable of introducing new hypothesis, it provides great heuristic values and makes the discovery straightforward (Cellucci, 2013b).

Similarity measures (e.g., Cosine similarity and Euclidean distance) are the manifestation of (dis)similarity pattern in the conceptual space. These measures introduce criteria to quantify semantic (dis)similarity. The similarity and distance measures have been vastly used in the field of cognition to model human categorization (Derrac & Schockaert, 2015).

3.5 Semantics underlying plausible patterns

Hierarchical-driven plausible patterns, like *generalization* and *specialization*, explore hierarchical relationships of the concepts in conventional ontological constructs (i.e., parent-child relationship) to draw new inferences. Ordered-based patterns conduct plausible inferences based on conceptual betweenness and partial order of concepts, definitions, actions, and phenomena. *Interpolation* and *a fortiori*, two ordered-based patterns, leverage measurable properties to compare and sort concepts regarding their size, order, location, ranking, etc. and infer new pieces of knowledge.

Equivalence plausible patterns, *similarity* and *dissimilarity*, draw new association between any comparable concepts using either hierarchical or ordered relationships, which imply any sort of equality (or inequality). However, equivalence patterns introduce two challenges: (i) they are highly context-dependent; two concepts that are similar in one aspect, might be totally different in other aspects; (ii) they are more beneficial when there are more of (dis)similar concepts (Derrac & Schockaert, 2015).

Although the mechanism of *metaphor* may seem comparable to *similarity* (i.e., they both search for commonalities between entities), the identification of the pre-existing similarities between two entities in metaphor relies on the cognitive meaning of the concepts, phrases, etc. in two different domains (the source and the target domains). This similarity is different than the epistemic commonalities between two concepts in one domain, which the *similarity* pattern looks for. Despite the conceptual similarities, which are intrinsic and static, the similarities in metaphor are (i) very subjective, (ii) hard to describe formally and challenging to capture and represent, and (iii) not necessarily symmetric (i.e., "Heart is made of glass" is a metaphor, but "Glass is made of heart" does not mean anything.), and (iv) significantly different from one domain to another. Likewise, the associations conducting *metonymy* (i) are contextual (e.g., wall street and stock exchange, or white house and president of the USA) and (ii) should have been observed prevalently in the past to make sense and be comprehensible for audience with no complication (Barcelona, 2012; Cellucci, 2013b).

Table 4 classifies the plausible patterns to four main categories based on the semantic relationships that they are applied to: (i) *generalization* and *specialization* exploit *hierarchical* relationships, (ii) *interpolation* and *a fortiori* are built upon *(partial) order* between concepts, (iii) *similarity* and *dissimilarity* leverages any association that imply some sort of equality (or inequality), and (iv) *metaphor* and *metonymy* rely on *cognitive meaning* and *mental perception* of concepts.

Table 4- Definition and classification of the plausible patterns based on the semantics they leverage

Type	Plausible Pattern	Description
Hierarchy-based patterns	Generalization	Passing from a given set of objects to a larger set that contains the given set.
	Specialization	Passing from a given set of objects to a smaller set that is contained in the given one.
Ordered-based patterns	Interpolation	Creating a new relationship from the observation space X to the conclusion space Y, where $x_i \in X$ is not mapped to any y $\in Y$ (unknown relationship), but other relationships from $x_h, x_j(\neq x_i)$ to Y and $x_h < x_i < x_j$ are known.
	A Fortiori	An inference from a proposition with high degree of confidence to a less confident proposition that is not clearly specified but is implicit in the first one.
Equivalent patterns	Similarity/ Dissimilarity	Moving between any two comparable nodes (siblings) in the concept hierarchy or comapring two overlapping (or non-overlapping) concepts w.r.t. a measurabel property.
Cognitive patterns	Methaphor	Considering that concepts of one domain belong to another domain, implying that if the concepts of the source domain have a specific property, the concepts from the other domain have the same property as well.
	Metonymy	Letting one concept stands for another concept, which is relevant and represents well-known characteritstics of the source conpcet.

3.5.1 Feasibility and challenges of implementing plausible patterns

The discussion above suggests that the feasibility of implementing the plausible patterns depends on two key criteria: (i) capturing and representing the underlying semantics that the patterns are applied to (i.e., knowledge representation), and (ii) identifying and formalizing the rationale behind the patterns (i.e., reasoning).

Description Logics (DLs) and related formalisms, including the ontology languages for the Semantic Web like OWL, offer well-defined semantics to represent and reason with taxonomic hierarchies. Description logics express definitions of *classes* and their *relationships* in a hierarchical structure. Class subsumption, and instance identification are two key capabilities of description logics that support automatic inference of class-subclass relationships through inheritance (Gil, 2005). Thus, DLs and, subsequently the SW ontologies, can support both knowledge representation and reasoning requirements in hierarchical patterns (i.e., generalization and specialization), and hierarchical equivalent patterns (i.e., similarity and dissimilarity inferences that leverage conceptual similarities).

However, description logics, and consequently OWL, do not essentially provide the required constructs to express sequential relationships with respect to a measurable property. While, the representation of the ordered relationships is a challenge, we identified the rationale of interpolation and a fortiori fathomable and practical to be formalized as a non-deductive form of logic. We assume the various reasoning paradigms offered by the Semantic Web (e.g., OWL reasoning, rule-based reasoning, query rewriting) can provide the required mechanisms to formalize and implement the logic behind the order-based patterns and generate new knowledge by analyzing the underlying ordered semantics.

Among the patterns above, we found the working mechanism of the cognitive patterns, *metaphor* and *metonymy*, different and sometimes intricate as they involve cognitive analysis. The semantics underlying cognitive patterns require mental counterparts for notions regarding a particular context. This prerequisite demands an access to mental spaces that are constructed from relationships between the concepts (aka., semantics) along with a perceptual experience in general and linguistic meaning in particular (Barcelona, 2012; Song, 2011). Hence, we found inducing symbolic and interpretable cognitive semantics and formalizing the cognitive plausible patterns a complicated task that imposes new challenges (e.g., linguistic analysis, sentiment analysis) that are beyond the scope of this work.

Our Semantic Web inspired approach to (i) represent *ordered* relationships among the *hierarchical* and *equivalent* associations within a unified ontology language and (ii) formalize and implement the *hierarchical*, *order-based* and *equivalent* patterns in one *plausible reasoner* is elaborated in the next chapter.

3.6 Implementations of plausible reasoning

There have been few attempts (Dontas & Zemankova, 1990; Oroumchian & Oddy, 1996) in different application domains to implement the theory of human plausible reasoning introduced by (Collins & Michalski, 1989). In their work, Collins and Michalski assumed that a larg body of human knowledge is stored and represented in hierarchies. By leveraging the hierarchcial constructs, they formalized four plausible patterns (generalizaiton, specializaiton, similairty and dissimilairty) that people use when they do not know a prompt answer. They also introduced a set of certainty factors (e.g., conditional likelihood, typicality) that calculates the certainty of the plausible answers. In their theory, the knowledge is represented in the form of logical statements (i.e., *descriptor(argument) = referent*). To perform the plausible patterns, they suggest transforming the argument or referent of the statement via plausible patterns, trying to find an (plausible) answer that satisfies the new statement.

Dontas et. al. (Dontas & Zemankova, 1990) developed a pilot version of the theory of human plausible reasoning (Collins & Michalski, 1989) on the periodic table. In their system (called APPLAUSE) they tried to generate the unknown (or deliberately deleted) attribute values of the elements using the known facts of other elements via the plausible patterns (including generalization, specialization, similarity and dissimilarity). To represent the domain knowledge, APPLAUSE utilizes the combination of hierarchies (nodes connected via parent-child relationships in periods and groups), statements (i.e., tuples in logic) describing entities in the hierarchy, similarity values between the entities, and dependencies and implications between properties of the entities. They used Prolog environment to develop APPLAUSE.

Oroumchian et. al. (Oroumchian & Oddy, 1996) investigated the applicablilty of human plausible reasoning (Collins & Michalski, 1989) in the domain of information retireval. They were motivated by the idea that people have been doing information retrieval long before the develpment of computers. In their study, they try to simulate the thinking process of librarirans when try to find relevant documents that may be interesting to a user. To represent the docuemtns, they used phrases and logical terms X and BN that occure in the phrases. Phrases are combination of two terms that are connceted (e.g., *User interface of Windows operating system*). An X realtionship expresses the relationship between a pharese and its parts (e.g., relationship between *32-bit* and *32-bit Operating System*) and a BN (Broader-Narrower) relationship represent the conenction between two following phrases (e.g., *relational* and *database*). The X and BN relationships are leveraged to demonstrate the hierarchical relationhsips between the pharases and later conduct the plausible reasoning via generalizaiton, specializaiton, or (dis)simialrity.

Although the experiments of both works show promising results, there are still some restrictions in the theory and implementation of the works that hinder the full potential of plausible reasoning, as an exploratory data analytic method.

The most important limitation of the theory, and correspondingly the systems above, is that they only consider hierarchical relationships in their inferences. As mentioned, the theory is founded on this assumption that large part of human knowledge is represented in hierarchies. Whereas, in real settings a variety of associations exist (e.g., cause/effect, action/product, partial order). Each of these associations carries different meaning and expresses interesting semantics between concepts that can be used in plausible reasoning and generate new associations.

In addition, the representation of knowledge in the form of statements impose some serious limitations. In a statement, an *argument* is associated with a *referent* via a *descriptor* (i.e., *descriptor(argument) = referent*). This schema does not allow the system to capture and represent complex relationships in real-world scenarios that may involve more than one argument. Furthermore, the statement representation of knowledge along with the statement transformation approach to implement plausible patterns compels the incoming questions to be formulated as statements. In the domains that questions are typically complicated and not simply possible to be stated in the form of a simple statement (such as questions with conditions in their WHERE clauses) the systems above will be ineffective. Hence, the statement representation of knowledge limits both the knowledge types that the system can contain and the queries that it can answer. A flexible and dynamic knowledge representation that efficiently captures and represents various forms of semantic associations offers a great potential to increase the scope of plausible reasoning.

Regardelss of the limtiations that the theory implies, there are some concerns regaridng the developed systems that renders their capabilites for real-world applications. The first concern is the size of the data and the scale of the experiments. APPLAUSE performed plausbile reasoning on the periodic table with 102 elements (at that time and 118 elements now) with shallow hierachies (i.e., two-level deep hierachies), and (Oroumchian & Oddy, 1996) evaluated their plausible information retrieval system in CACM collection that only includes 3,204 documents. Despite the fact that both the systems present a *concept of proof* implementation of the theory, neither of them report any perfomance results. So, the applicabilty of their implementations in applications with large data is questionable.

In addition, despite the emphasis of the theory on the significance of domian knowledge in conducting plausible reasoning, no standard domain knowledge is utilized in the plausible information retrieval system. Only simple clue-based methods were used to retrieve the relationships and construct the knowledge base. Their further investigation showed that 90% of the relationships were accurate and many relationships were not identified (Oroumchian & Oddy, 1996).

Amongst the analytics systems that implement plausible reasoning to derive new knowledge, there are some studies that leverage the concept of *plausibility* to introduce novel approaches to handle uncertainty. For example, Schechter (Schechter, 2015) combined the logic of plausibility with the logic of justification to introduce a *logic of plausible justification*, which used to develop an argumentation and debate platform for multi-agent systems. The notion of plausibility allows the agents to hold a incorrect and unreliable belief, as long as they have a plausible evidence (or justification) for it. In their system, agents can discuss the correctness of a belief and the plausibility of the evidence.

The mechanism and definition of plausible reasoning has not always been unambiguous that clearly addresses the notion of plausibility (introduced early in this chapter). In some studies, plausible reasoning has been mistakenly used to imply any kind of approximate reasoning in complex and uncertain settings, including probability (Horvitz, Heckerman, & Langlotz, 1986; Prade, 1985), fuzzy logic (Bouchon-Meunier, Dubois, Godo, & Prade, 1999), and Dempster-Shafer theory (Dezert, 2002).

3.7 Summary

This chapter presented plausible reasoning as a non-deductive, data-driven reasoning approach exploring semantic data to infer new knowledge supported by reasonable evidence. We studied eight patterns (including generalization, specialization, similarity, dissimilarity, a fortiori, interpolation, metaphor, metonymy) that navigate the relationships between semantic data (including hierarchical, ordered, equivalent and cognitive) to draw new inferences. The plausible inferences are non-demonstrative, ampliative, non-monotonic and subjective. Despite the recognized challenges with formalizing the hierarchical, order-based, and equivalent plausible patterns and representing ordered-based relationships, we found the cognitive patterns irrelevant to the objectives of this study (at least at this current step of the work).

This chapter showed, although a body of theoretical work in different disciplines (e.g., philosophy, mathematics, artificial intelligence) has been done to explain and formalize plausible reasoning, its potential to analyze semantically-represented data and discover interesting relationships has not been investigated in practice. Given the availability of a variety of sources of data (i.e., big data), along with novel knowledge representation and reasoning approaches designed for semantic analysis and discovery (i.e., the Semantic Web technologies) bring to attention the need to investigate the potential of plausible reasoning in big data analytics.

Chapter 4: Plausible Reasoning Over Knowledge Graphs

After discussing the fundamental concepts of this thesis in the previous chapters, we are now able to provide a formal definition of our plausible reasoning approach by identifying its components, the plausible patterns and the semantics that they are based upon. These definitions clarify the technical aspects of our approach to implementing plausible reasoning over knowledge graphs. Later, the challenges of the implementation, the potential solutions and our contributions to address them are discussed. This is then followed by the technical details of our solutions and how they were applied to implement a plausible reasoner. Finally, we introduce the SeDan framework, as a pragmatic endeavour to achieve plausible reasoning in a real-world setting.

4.1 Definitions

Plausible reasoning has been used in a wide range of studies and its definition, which has not always been unambiguous, has evolved over time. Along with the objectives of this thesis, we provide a formal definition of plausible reasoning conforming with the Semantic Web inspired approach that we have taken. We formulate a series of definitions formalizing the concepts that plausible reasoning is built upon. Leveraging those components, we ultimately define plausible reasoning over knowledge graphs.

Definition 1 (Knowledge Graph[2]): a knowledge graph is a directed labeled multigraph where entities/concepts are represented by nodes, and edges represent their relations. In a knowledge graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{R}, \mathcal{O})$, $\mathcal{V}$ is the set of entities, $\mathcal{E}$ is a set of labeled directed

[2] Knowledge graph has been widely applied in different applications, while its definition has not always been consistently agreed upon. Thus, including a definition that serves as a basis for our discussion and provides a common vision of knowledge graph is important.

edges between two entities, $\mathcal{R}$ is the set of the predicate labels and $\mathcal{O}$ is the ontology of the entities. In a knowledge graph, the triple $\langle s, p, o \rangle$ represents a directed relation $p \in \mathcal{E}$ from a subject entity s to the object entity o $(s, o \in \mathcal{V})$ (Das et al., 2017; B. Shi & Weninger, 2016).

Definition 2 (Semantic Association): a semantic association, $\sigma \in \Sigma$, implies a meaningful relationship between two concepts. In a graph representation of knowledge, a semantic association is equal to a directed labeled link connecting two nodes $(s, o \in \mathcal{V})$ via an edge $(p \in \mathcal{E})$. A semantic association can indicate super/sub-ordination relationships via hierarchical constructs (i.e., $\in$ and $\subset$), equality relationships via equivalence constructs (i.e., $=$ and $\neq$), or interpretable directions via sequential associations (i.e., $\prec$ and $\succ$).

Definition 3 (Plausible Pattern): a plausible pattern, $\pi \in \Pi$, provides the rational criteria for exploring the knowledge graph by navigating from one entity (node) to another entity (node) based on a set of well-defined semantic relationships. In the development of plausible reasoning over a knowledge graph, $\mathcal{G}$, plausible patterns can be considered as pattern matching functions, $\phi(\nu, \pi, \mathcal{G})$, that identify the triples, $\langle \nu, \sigma, \nu' \rangle$, acknowledging the logic of the corresponding plausible pattern, π—i.e., two concepts, ν and ν' are connected via a semantic association, σ, conforming with logic of the plausible pattern, π:

$$\phi(\nu, \pi, \mathcal{G}) = \{\nu' | \langle \nu, \sigma, \nu' \rangle, \sigma \in \Sigma, \sigma \Longrightarrow \pi\}$$

in which $\nu, \nu' \in \mathcal{V}$, $\Sigma = \{\in, \subset, \prec, \succ, =, \neq, \ldots\}$, and $\Pi = \{GEN,\ SPEC,\ INTP,\ AFORT,\ SIM,\ DIS\}$. Table 5 presents the variants of the pattern matching function using relevant associations to retrieve the applicable semantics to each plausible pattern.

Table 5- The pattern matching functions based on the plausible patterns applied to semantic associations-there are two variants of the matching function for each of interpolation and a fortiori, as they can be applied in two directions. Similarity and dissimilarity matching functions include two patterns, since the equivalent (or inequivalent) associations are symmetric.

Plausible Pattern	Semantic Associations	Pattern Matching Function
Generalization	$\Sigma = \{\in, \subset\}$	$\phi(v, GEN, \mathcal{G}) = \{v' \mid \langle v, \sigma, v' \rangle, \sigma \in \Sigma\}$
Specialization	$\Sigma = \{\in, \subset\}$	$\phi(v, SPEC, \mathcal{G}) = \{v' \mid \langle v', \sigma, v \rangle, \sigma \in \Sigma\}$
Interpolation	$\Sigma_1 = \{<, \prec\}$ $\Sigma_2 = \{>, \succ\}$	$\phi(v, INTP, \mathcal{G}) = \{v', v'' \mid \langle v', \sigma, v \rangle \langle v, \sigma, v'' \rangle, \sigma \in \Sigma_1\}$ $\phi(v, INTP, \mathcal{G}) = \{v', v'' \mid \langle v'', \sigma, v \rangle \langle v, \sigma, v' \rangle, \sigma \in \Sigma_2\}$
A Fortiori	$\Sigma_1 = \{<, \prec\}$ $\Sigma_2 = \{>, \succ\}$	$\phi(v, AFORT-L2M, \mathcal{G}) = \{v' \mid \langle v, \sigma, v' \rangle, \sigma \in \Sigma_1\}$ $\phi(v, AFORT-M2L, \mathcal{G}) = \{v' \mid \langle v, \sigma, v' \rangle, \sigma \in \Sigma_2\}$
Similarity	$\Sigma = \{=, \approx, \equiv, \sim\}$	$\phi(v, SIM, \mathcal{G}) = \{v' \mid \langle v, \sigma, v' \rangle \ OR \ \langle v', \sigma, v \rangle, \sigma \in \Sigma_1\}$
Dissimilarity	$\Sigma = \{\neq, \not\approx, \not\equiv, \nsim\}$	$\phi(v, DIS, \mathcal{G}) = \{v' \mid \langle v, \sigma, v' \rangle \ OR \ \langle v', \sigma, v \rangle, \sigma \in \Sigma_1\}$

Definition 4 (Plausible Path): the path between two nodes in a knowledge graph connected via a (series of) semantic association(s) conforming with the logic of plausible patterns—i.e., a chain of semantics retrieved by pattern matching functions. A plausible path connects a source node to a target node via a single or a sequence of semantic associations applicable to plausible patterns through a set of intermediate nodes:

$$plPath = \{\langle v, \sigma, v_1 \rangle \langle v_1, \sigma_1, v_2 \rangle \ldots \langle v_n, \sigma_n, v' \rangle \mid \pi_n \in \Pi, n \geq 0\}$$

Alternatively, we can define a plausible path as a sequence of plausible patterns:

$$plPath = (\pi_1, \ldots, \pi_n) \, ; \quad \pi_n = \langle v_{n-1}, \sigma_n, v_n \rangle, n \geq 1$$

Definition 5 (Plausible Association): a new association, $\langle v, \varepsilon, v' \rangle$, connecting the source node to the target node (or vice versa) of a plausible path.

Definition 6 (Plausible Reasoning): in a ***knowledge graph***, ***plausible reasoning*** is the act of deriving a new line of inference (i.e., ***plausible association***) by connecting two previously disconnected nodes via a ***plausible path*** (i.e., comprised of ***plausible patterns***), whose consistency can be further evaluated deductively and may lead to the discovery of new *solutions* whose reliability is *plausible* (Figure 2)

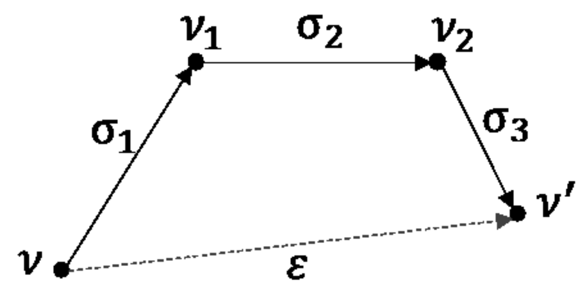

Figure 2- Plausible reasoning over knowledge graphs - $\langle \nu, \sigma_1, \nu_1 \rangle \langle \nu_1, \sigma_2, \nu_2 \rangle \langle \nu_2, \sigma_3, \nu' \rangle$ *is a plausible path comprised of 3 plausible patterns. In the plausible triple,* $\langle \nu, \varepsilon, \nu' \rangle$*,* ε *is the plausible association.*

Let $\mathcal{G}$ be a knowledge graph including ontological constructs $\mathcal{T}$ and (incomplete) assertional data $\mathcal{A}$ ($\mathcal{G} = \langle \mathcal{T}, \mathcal{A} \rangle$) and Q a query representable in the triple format $\langle s, p, o \rangle$, a plausible reasoner $plRes(\mathcal{G}, \mathcal{T}, \mathcal{A})$ returns a set of solutions for Q:

$$\mathrm{P}^{\mathcal{G}}_{Q=(s,p,o)} = \{(plAns_i \, , plPath_i) | \; i \; \geq 0\}$$

$plAns$ is a plausibly inferred solution and $plPath$ expresses the corresponding chain of semantic relationship(s) leveraged by the corresponding plausible patterns. In the formula above, the index of plausible answer starts from 0 ($i \; \geq 0$), since a plausible path does not guarantee an answer.

The act of plausible reasoning: our definition of plausible reasoning (definition 6) suggests inferring a plausible answer could be formulated as the problem of deriving a plausible association (definition 5) interrelating two unconnected entities via a plausible path (definition 4), which is comprised of a single or a sequence of plausible patterns (definition 3). The act of plausible reasoning over knowledge graphs can be formulated as a four-stage process:

1. Receive a question to answer and the already available data (i.e., *what we look for*),

2. Utilize plausible patterns to conduct an exploratory search over the data and collect supporting semantic relationships relevant to the concept(s) in the question,
3. Draw a new association by reformulating the question (e.g., making hypothesis in scientific discovery) using the semantics retrieved in stage *2*: replacement of the target entities with the entity in the question will generate plausible questions that can be tested deductively (i.e., *how we look for*),
4. Evaluate the plausibly generated questions (i.e., hypotheses) and interpret the results and communicate the plausible answers (i.e., *how we interpret*).

In practice, a plausible reasoner, like any analytic method in scientific discoveries, starts with a problem (i.e., expressed in the form of a question) and the already available data to make new hypotheses that can be tested further (Cellucci, 2013b; Collins & Michalski, 1989).

4.2 Requirements and challenges of implementing plausible reasoning

According to the definition of plausible reasoning provided earlier, developing a plausible reasoner has two main aspects: the representation of data/knowledge and the reasoning paradigm. The first aspect should address how a knowledge representation formalism can offer the required expressivity to capture real-world entities and their properties, while carrying their semantics. The second aspect tries to find a reasoning approach that reinforces the analysis of semantic data while conforming with the rationale underlying the plausible patterns. The following sections elaborate the requirements of each aspects, the challenges that they pose and the solutions to address those challenges.

4.2.1 Knowledge Representation

The definition of plausible semantics (definition 2) implies that, in addition to the real-word entities and their attributes, the relationships between those entities are of importance. The semantics of an entity is identified via a clear and unambiguous relationship with other entities. Furthermore, the definition of plausible reasoning (definition 6) indicates that the plausible reasoner leverages both data (i.e., assertional data) and domain knowledge (i.e., ontological constructs) to perform a chain of inferences and reach a plausible answer. Hence, the *connectivity* of data items within a source and to the relevant, external sources is paramount.

Definition 5 (Plausible Association) suggests that the knowledge representation formalism should be capable of including *arbitrary* relationships that are drawn within the process of plausible reasoning and ultimately stores them as plausible answers. Additionally, the fourth step of the *plausible reasoning in act* raises the issue of data *provenance*, in which a plausible answer is justified. To provide users with reliable answers, the evaluation of the plausible answers traces the reasoning process and explains where the knowledge came from.

Hence, an appropriate knowledge representation formalism for an effective plausible reasoning should be *connected* (i.e., entities are interrelated), *expressive* (i.e., carries semantics), *flexible* (i.e., includes relationships that may not be restricted in their domain and/or range) and encode the *provenance* (i.e., justifies where the knowledge came from).

As discussed in chapter 2, knowledge graph offers an ideal knowledge representation formalism mitigating the requirements above. The Semantic Web technologies fundamentally provide ideal technical means reinforcing the development of graph-based representation of knowledge. The Semantic Web offers a graph-based representation of data via RDF, and RDFS and OWL, which carry the semantics of data more profoundly.

Resource Description Framework (RDF) serializes knowledge/data points via a *connected* series of triples in the form of subject-predicate-object. The RDF design allows the creation of new *arbitrary* graph constructs (e.g., new relations, or chain of connections), while providing the fundamentals to build more *expressive* ontology languages like RDFS and OWL, which can associate specific *semantics* within the data graph. Linked data interlinks different sources of data/knowledge as one connected *large* graph, while their origin is still *trackable* (Paulheim, 2017; Van Ossenbruggen, Nack, & Hardman, 2004).

Knowledge graphs in the Semantic Web framework provides the means to represent and subsequently analyze the underlying semantics of the available data via plausible reasoning. However, even OWL, the most expressive knowledge representation language in the Semantic Web framework, does not fully support all the semantics that the plausible patterns are applied to.

The built-in RDFS constructs (like *subClassOf*, *subPropertyOf*) and OWL axioms (like *instanceOf*, *sameAs*, *differentFrom*, *equivalentClass*) offer the necessary semantics to manifest the hierarchical (i.e., generalization and specialization) and hierarchical-equivalence (i.e., hierarchical similarity and dissimilarity) plausible patterns (Table 5). The ordered-based (i.e., interpolation and a fortiori) and ordered-equivalence (i.e., ordered similarity and dissimilarity) plausible patterns conduct the reasoning based on measurable relationships between comparable concepts, while the existing relationships in RDFS or OWL cannot represent and reason with them effectively.

Hence, the lack of expressivity of the Semantic Web ontology languages raises the first challenge: to be able to implement plausible reasoning within the logic layer of the Semantic Web, we need to introduce an augmentation to OWL that fulfills the representation and, subsequently, reasoning with *all* the semantics underlying the plausible patterns in one integrated/solid framework.

4.2.2 Reasoning paradigm

As discussed in chapter 2, the Semantic Web technologies offer a variety of reasoning mechanisms that each one is suitable for different applications and domains, regarding the size of the data, complexity of the ontology, and the reasoning requirements (i.e., handling uncertainty). Now the question is which reasoning approach the most appropriate mechanism is to implement plausible reasoning over large data represented in knowledge graphs.

The Semantic Web ontology languages (e.g., RDFS and OWL) offer the constructs to semantically markup data items in the Semantic Web, but they are based on formal logic (e.g., description logics) capturing general constraints that are globally accepted (i.e., rules). For the very same reason, automated reasoning via these logic-based ontology languages (i.e., OWL 2 EL) on RDF graphs is constrained to deductive reasoning tasks inferring accurate, sound and consistent (i.e., with no logical contradictions) relationships among the classes, properties and relationships (World Wide Web Consortium, 2012). Hence, conventional ontology-based reasoning lacks the advantages of non-demonstrative, explorative reasonings (e.g., plausible reasoning) due to the expressivity limitations residing in the deterministic knowledge representation languages.

Leveraging the ontology languages (e.g., OWL), rule languages (e.g., SWRL) are proposed to improve the expressivity of the Semantic Web. However, despite the efficiency of rule-based approaches, there are some challenges in exploiting rule-based reasoning that hinder an effective implementation of plausible reasoning over large datasets. Traditional rule engines suffer from scalability. Performing rule inferences (e.g., OWL 2 RL or SWRL) on large datasets can take hours to finish. Even in the case of multi-core and multi-CPU machines, rule-based inferences should be efficiently parallelized. Although recent efforts have proposed distributed rule engines to handle this challenge, but performance of rule-based reasoning with regard to large knowledge bases is still an issue (Kolovski, Wu, & Eadon, 2010; J. Zhang, Yang, & Li, 2017). Additionally, it is not always possible to capture and express the domain knowledge in the form of rules. More importantly, rule-based reasoning is the process of applying a set of rules to a set of statements (e.g., triples) to return some conclusions (i.e., answer set) (Urbani, Van Harmelen, Schlobach, & Bal, 2011). Hence, the deductive rule reasonings materialize what has been already captured in the domain knowledge in the form of rules—i.e., rule-based reasoning has no ampliative power enabling assumption seeking or generating hypotheses.

The latest version of the SPARQL query language, SPARQL 1.1, supports the use of ontological constructs to augment query answering over RDF graphs under logical entailments (i.e., entailment regime). For this purpose, OWL 2 QL fragment offers a dedicated schema to rewrite SPARQL queries (i.e., into the form of a conjunctive query) that will be further evaluated over RDF data. Although exploiting ontological constructs to rewrite a query has been in practice for decades (i.e., query rewriting has been one of the approaches to ontology-based data access), the new features of SPARQL 1.1, including property paths, value creation, etc. (Harris & Seaborne, 2013) provide this innovative opportunity to rewrite a query using both the ontological part of the knowledge and the assertional data. Especially that the knowledge graphs, as a modern formalism for representing knowledge, effectively combine variety of knowledge and data sources in one integrated platform (Bischof, Krötzsch, Polleres, & Rudolph, 2014; Glimm & Ogbuji, 2013).

Thus, while OWL 2 QL supports query rewriting to preform standard reasoning tasks, SPARQL 1.1 reinforces more complex reasoning paradigms incorporating data-driven approaches. Hence, query rewriting over OWL 2 QL using SPARQL 1.1 can be considered as a flexible reasoning approach that can infer logical, deductive answers, while it is capable of leveraging data associations (i.e., semantics) to derive new plausible solutions.

However, as discussed before, a query rewriting technique requires a set of conditions (aka. inference constraints) to conduct the rewriting and infer new solutions. Hence, the challenge of implementing plausible reasoning via query rewriting over OWL QL using SPARQL 1.1 could be narrowed down to the identification of the rationale of plausible patterns and formalizing them in the form of inference constraints conducting the query rewriting. In the next section we discuss how the formalization of the plausible patterns (definition 3) via the pattern matching function (Table 5) tackles this challenge.

4.3 Solutions

This research aims to investigate the potential of implementing plausible reasoning over knowledge graphs, targeting a semantic analytics framework for large health data analytics. We aim to demonstrate the effectiveness of non-deductive exploratory analytics within the logic layer of the Semantic Web to infer new knowledge from large data-sets, especially when we are dealing with incomplete and noisy data.

However, as discussed, despite the efficient means that Semantic Web technologies offer to analyze the underlying semantics of the available knowledge and facilitate the implementation of plausible reasoning over knowledge graphs, they are still two main shortcomings: the limited expressivity and reasoning obstacles. To address these challenges, we show:

i. The lack of expressivity of the Semantic Web ontology languages, can be addressed by introducing additional markups (we call it PLausible OWL extension, PL-OWL) that extend OWL constructs by providing a greater flexibility in modeling and representing the semantics applied to the plausible patterns.

ii. To implement plausible reasoning via plausible patterns, query rewriting is an appropriate reasoning paradigm on the Semantic Web that leverages OWL constraints (can be extended by our PL-OWL) to transform a given SPARQL query allowing the extraction of both explicit (deductive) and implicit (plausible) knowledge from the underlying data. Formalizing and implementing the rationale behind the plausible patterns in the form of a set of inference constraints residing in the core of the query rewriting algorithm can deliver plausible reasoning. And ultimately,

iii. The integration of the contributions as a SEmantics-based Data ANalytics (SeDan) framework establishes the *act of plausible reasoning* (definition 6) in a working system to discover new associations between underlying domain-specific data when deductive query answering fails. The framework is evaluated using health data, since semantic analytics is found relevant to healthcare (Mohammadhassanzadeh, Woensel, Abidi, & Abidi, 2016), as a predominantly knowledge-intensive domain supporting both diagnosis reasoning and predictive inference.

The sections below elaborate the theoretical and technical aspects of each of the contributions that this work offers.

4.3.1 Plausible OWL extension

Within the Semantic Web framework, standard reasoning capabilities of OWL profiles (such as OWL 2 QL) support various types of ontology-based inference. Subsumption (i.e., class) relationship properties (i.e., *rdfs:subClassOf* and *owl:instanceOf*) support generalization and specialization patterns by moving between the nodes in a hierarchical (taxonomic) structure—i.e., from parent to child or vice versa. Equality semantics (i.e, *owl:sameAs and owl:disjointFrom*) conduct the equivalence patterns by moving between similar and interchangeable nodes in the ontology.

However, OWL does not support all the semantics required in the plausible patterns. The order-based patterns, a fortiori and interpolation, conduct plausible reasoning based on measurable relationships (e.g., size, chronological order, location, ranking, phase, etc.) between comparable concepts, objects or actions. The existing constructs in RDF(S) or OWL cannot represent and reason over the measurable relationships effectively. Hence, we (Mohammadhassanzadeh, Raza Abidi, Shah, Karamollahi, & Abidi, 2017) have introduced a plausible OWL extension to represent order-based semantics within the SW framework.

Despite the probabilistic/fuzzy extensions to OWL (R. N. Carvalho et al., 2017; P. da Costa et al., 2008; Dong et al., 2015; Liu et al., 2013) that incorporate new type of individuals or uncertainty/belief/truth values describing facts, the plausible OWL extension (PL-OWL) does not introduce any new types of facts or values about entities. PL-OWL includes a set of classes, sub-classes, properties and sub-properties enabling OWL ontologies to incorporate (i) ordered relationships and (ii) new notions introduced in the theory of plausible reasoning (e.g., plausible pattern, plausibly inferred answer). The plausible augmentation to OWL mainly includes a new type of property (and its sub-properties) defining the *ordered* interrelations and restrictions. These properties express new types of semantics in the form of associations between entities. To be consistent with the OWL axioms, we also need to introduce some new classes characterizing the ordered properties.

As a subclass of owl:ObjectProperty, the class *OrderedProperty* supplements the constructs of OWL to capture ordered relationships and support all three types of semantic associations in one coherent ontology language (definition 2). The two subclasses of OrderedProperty, *StandsAfter* and *StandsBefore*, provide extra expressivity identifying the direction of the comparison being made in interpolation and a fortiori (Table 5).

The class of *PlausiblePattern* and its individuals (e.g., generalization, interpolation, etc.) represent the plausible patterns (definition 3) that are formalized in the form of inference constraints and integrated into the process of plausible reasoning (i.e., conduct the query rewriting). The class of *Context* captures the domain specific knowledge with respect to specific conditions to address ambiguity, especially in the case of ordered relationships where the associations between the items are highly dependent on the context in which there are being compared.

Subsequently, there are some extension constructs supplementing the PL-OWL properties. The properties *standsAfter* and *standsBefore* provide the opportunity to represent the ordered relationships that the domain ontology has no label for them. The instances of the class *Context* help to append the circumstances to the ordered property. But, if there is any labeled associations (e.g., precedes, bigger) implying an ordered relationship, it can be directly assigned to the *StandsAfter* and *StandsBefore* classes (this process will be explained further in the plausible OWL enrichment section).

Table 6 demonstrates the proposed OWL extension followed by detailed description of each construct.

Table 6- Plausible OWL extension (PLOWL)

Class Name	Supper Class		On Property	
OrderedProperty	ObjectProperty		-	
StandsAfter	OrderedProperty		standsAfter	
StandsBefore	OrderedProperty		standsBefore	
Context	Class		hasContext	
PlausiblePattern	Class		inferredViaPattern	
PlausibleAnswer	Class		-	
Property Name	**Type**	**Domain**	**Range**	**Inverse Property**
standsAfter	StandsAfter	Entity	Entity	standsBefore
standsBefore	StandsBefore	Entity	Entity	standsAfter
hasContext	ObjectProperty	Entity	Context	-
inferredViaPattern	ObjectProperty	PlausibleAnswer	PlausiblePattern	-

plowl:OrderedProperty is a class of properties that represents partial order between two classes or entities w.r.t a measurable property (i.e., *plowl:Context*). More formally, if P is an *OrderedProperty*, any instance of P, like $(x\ P\ y)$, implies a sequence or a relative quantity between x and y – i.e., x is bigger, older, slower, etc. than y or vice versa. From this, the plausible reasoner would be able to conduct *interpolation* and *a fortiori* reasoning. *plowl:StandsAfter* and *plowl:StandsBefore* are sub-classes of *plowl:OrderedProperty*, identifying the direction of a sequence.

Ordered Property

Description

The class of ordered properties. Like owl:TransitiveProperty, owl:FunctionalProperty, etc., having a separate class for representing ordered properties makes the modeling of and reasoning with plausible semantics easier.

Properties with OrderedProperty as their domain or range

-

Instances of OrderedProperty

StandsAfter

StandsBefore

plowl:StandsAfter is defined as a class of properties that demonstrate the relationship between two classes, in which the first class *stands after* the second class ($\langle c_1\ StandsAfter\ c_2\rangle = \langle c_1 \succ c_1\rangle$). *Older*, *taller* and *succeeds* can be classified as *StandsAfter* relationships.

plowl:StandsBefore is defined as a class of properties that demonstrate the relationship between two classes, in which the first class *stands before* the second class ($\langle c_1\ StandsBefore\ c_2\rangle = \langle c_1 \prec c_2\rangle$). *Younger*, *shorter* and *precedes* can be classified as *StandsBefore* relationships.

StandsAfter

Description

Represents a class of properties that link any two comparable objects and expresses the (partial) order of them regarding a specific context.

Properties with StandsAfter as their domain or range

-

Instances of OrderedProperty

standsAfter

StandsBefore

Description

Represents a class of properties that link any two comparable objects and expresses the (partial) order of them regarding a specific context.

Properties with OrderedProperty as their domain or range

-

Instances of OrderedProperty

standsBefore

plowl:Context identifies the setting in which two concepts are being compared. Some ordered-based properties intrinsically imply a context – i.e., *older, shorter and precedes* imply *age*, *height* and *time* respectively. However, there are some properties that indicate a general *sequence* – e.g., *higher*. Hence, *Context* help to clearly distinguish the conditions that a generic order property is representing. For example, the statement *"the cancer risk of warts is higher than the herpes,"* can be represented as *higher(warts, herpes, cancer_risk)*, instead of defining the new property *higherCancerRisk*.

In addition, *Context* helps to implement the plausible OWL extension more efficiently. Without *Context*, the ontology representing the domain knowledge requires one different property to represent each generic ordered property, which drastically affects the complexity of the OWL ontology. While, considering *Context* in the extension helps to represent a generic order property via a pair of ordered-property and a context, e.g., *(standsAfter, cancer_risk)*.

Context

Description

Context is a constraint for a triple statement to disambiguate its meaning and express the circumstances in which the statement makes sense.

Properties with Context as their domain or range

hasContext (range)
inTheContextOf (range)

Instances of Context

will be determined by ontology engineer, based on the application domain

plowl:PlausiblePattern is a class representing the plausible pattern flavors implemented in the plausible reasoner. Currently, *PlausiblePattern* is a simple class with 6 instances: *generalization*, *specialization*, *similarity*, *dissimilarity*, *a fortiori*, and *interpolation*.

Plausible Patterns

Description

Includes the six well-known plausible patterns that we are implementing in this research.

Properties with PlPattern as their domain or range

inferredThroughPattern (range)

Instances of PlPattern

Generalization, Specialization, (Dis)Similarity, A fortiori, Interpolation

plowl:PlausibleAnswer is a class to represent the plausibly inferred answers, which are simply in the form of a triple statement. A plausible answer may be accompanied (*plowl:inferredViaPattern*) by the plausible pattern(s) (*plowl:PlausiblePattern*) that conducted the inference.

Plausible Answer

Description

Represents a plausibly inferred association, including a triple statement, the context in which the triple is tangible, and the pattern that has led to this inference.

Properties with PlausibleAnswer as their domain or range

inferredViaPattern (domain)
hasContext (domain)

Instances of PlausibleAnswer

Will be inferred via plausible reasoning using assertional knowledge

plowl:hasContext – in the cases that the RDF repository is not able to represent a context-based relationship in the form of a quadruple –i.e., *predicate(subject, object, context)*– *hasContext* will link a triple to its corresponding context, *hasContext(URI(predicate(subject, object)), context)*.

hasContext

Type: ObjectProperty

Description:

This object property links a property (predicate) to a context to disambiguate the meaning of a triple statement.

plowl:standsAfter is an instance of the *StandsAfter* class, representing a generic order property w.r.t a context. This property is an endeavor to represent and formalize ordered-based relationships between two entities/classes that no equivalent ordered-based property (*plowl:OrderedProperty*) represents them in the domain ontology. For example, the statement *higher(warts, herpes, cancer_risk)* can be formalized as *standsAfter(warts, herpes, cancer_risk)*. Having a standard axiom for representing the ordered-based properties reduces the complexity of the domain ontology and facilitates the implementation of the plausible patterns.

plowl:standsBefore is an instance of *StandsBefore* class (i.e., similar to *plowl:standsAfter*, but in the reverse order) representing a generic order property w.r.t a context. For example, the statement *higher(warts, herpes, cancer_risk)* can be formalized as *standsBefore(herpes, warts, cancer_risk)*.

standsAfter
Type: OrderedProperty (StandsAfter)
Description:
This object property is a link between any two comparable objects and expresses the (partial) order of them regarding a specific context. standsAfter(a,b,(c)) expresses that object *a* locates after object *b* regarding the context c. Context c can imply size, volume, ranking, level, etc. For example, if *c = length*, then a is longer than b.

standsBefore
Type: OrderedProperty (StandsBefore)
Description:
This object property is a link between any two comparable objects and expresses the (partial) order of them regarding a specific context. standsBefore(a,b,(c)) expresses object a locates before object b regarding the context c.

plowl:InferedViaPattern connects an inferred plausible answer to the pattern(s) that were used in the process of plausible reasoning.

inferredViaPattern
Type: ObjectProperty
Description:
This object property connects a plausibly inferred answer (PlAnswer) to the plausible pattern(s) that led to the inference.

Code 1 shows a snapshot of the implementation of the plausible OWL extension (complete PL-OWL Extension Code can be found in Appendix I). Based on these constructs, we will be able to enrich (Code 2) the existing OWL ontologies to represent ordered-based relationships (i.e., *SemMedDB:precedes a plwol:StandsBefore*), and support the plausible reasoner to conduct interpolation and a fortiori reasonings.

```
plowl:OrderedProperty a owl:Class ;
   rdfs:label "OrderedProperty" ;
   rdfs:comment "The class of (partial) ordered properties." ;
   rdfs:subClassOf owl:ObjectProperty.
plowl:StandsBefore a owl:OrderedProperty;
   rdfs:range owl:Thing;
   rdfs:domain owl:Thing;
   rdfs:comment "This object property is used to show which concept
   (subject) is located before another concept (object) regarding
   a specific context. The inverse property is
   StandsAfter."^^xsd:string.
plowl:hasContext a owl:ObjectProperty;
   rdfs:range plowl:Context;
   rdfs:comment "This object property links an object property to
   the context nodes being applied to it."^^xsd:string.
```

***Code 1**- Implementation of some of the constructs of the plausible OWL extension*

4.3.2 Plausible OWL enrichment

In addition to the plausible OWL extension, we introduced a plausible OWL enrichment to better distinguish, identify and formalize the plausible semantics underlying the plausible patterns. The PL-OWL enrichment works as an upper-level ontology providing a semantic foundation, which conducts the variants of hierarchical and equivalence patterns more effectively (Table 5). The plausible enrichment provides supplementary clarity facilitating the implementation of the plausible patterns in the query rewriting algorithm (the plausible query rewriting will be explained in the next section).

As Table 7 demonstrates, this upper-level ontology is created by merging constructs from OWL and PL-OWL. *HierarchicalProperty* represents any associations that expresses subsumption relationships. An instance of *HierarchicalProperty* can conduct generalizations or specialization. Likewise, *EquivalentProperty* expresses associations implying any type of similarity (or dissimilarity). A *SimilarityProperty* or *DissimilarityProperty* is further drilled down to the *ordered* and *hierarchical equivalent* properties, conducting ordered-equivalent patterns or hierarchical-equivalent patterns respectively.

Having the plausible OWL extension and plausible OWL enrichment implemented, now we need to enrich the existing data and domain ontology for two main purposes: (i) supplementing the domain ontology with ordered-based properties by providing the required semantics to conduct ordered-based plausible patterns, a fortiori and interpolation, and (ii) organizing the existing associations in the ontology into three main semantic associations with the aim of facilitating the implementation of the plausible patterns.

Table 7- Plausible OWL enrichment

Class Name	Supper Class	Subclasses / Instances
HierarchicalProperty	owl:ObjectProperty	rdf:type, db:substructure*
EquivalentProperty	owl:ObjectProperty	-
SimilarityProperty	plowl:EquivalentProperty	-
DissimilarityProperty	plowl:EquivalentProperty	-
OrderedSimProperty	plowl:SimilarityProperty	sem:occuresIn, sem:coexistWith**
HierarchicalSimProperty	plowl:SimilarityProperty	owl:sameAs, obo:hasExactSynonym***
OrderedDissimProperty	plowl:DissimilarityProperty	-
HierarchicalDissimProperty	plowl:DissimilarityProperty	owl:disjointWith, sem:differentFrom
* obo: Disease Ontology, ** sem: SemMedDB, *** db: Drug Bank		

Code 2 demonstrates some ontology enrichment (based on the constructs introduced in Table 6 and Table 7) of the knowledge sources used in the evaluation framework.

```
SemMedDB:PRECEDES          a   plwol:StandsBefore;
SemMedDB:lower_than        a   plwol:StandsBefore;
SemMedDB:CAUSES            a   plwol:StandsBefore;
db:substructure            a   plwol:HierarchicalProperty;
SemMedDB:ISA               a   plwol:HierarchicalProperty;
SemMedDB:PART_OF           a   plwol:HierarchicalProperty;
oboInOwl:hasExactSynonym   a   plowl:HierarchicalSimProperty;
SemMedDB:OCCURS_IN         a   plowl:OrderedSimProperty;
SemMedDB:COEXISTS_WITH     a   plowl:OrderedSimProperty;
SemMedDB:different_from    a   plwol:HierarchicalDissimProperty;
SemMedDB:different_than    a   plwol:HierarchicalDissimProperty;
```

Code 2- *Snapshot of the enriched ontology using the introduced plausible OWL extension*

4.3.3 Plausible query rewriting algorithm

As discussed earlier, we utilize query rewriting as a technique to implement plausible patterns and explore the knowledge graph with the aim of deriving new assertions (i.e., entailment) answering a failed query that was initially unresolvable. In our plausible query rewriting algorithm, we adopted the general idea of CGLLR algorithm (Calvanese, Giacomo, & Lembo, 2007), which explores the domain knowledge in an iterative process looking for applicable semantics to the body atom of the query. To retrieve the applicable semantics, the CGLLR algorithm utilizes a set of rewriting rules based on a partial function over a $DL - Lite$ ontology.

Our plausible query rewriting algorithm (Mohammadhassanzadeh, Abidi, Van Woensel, & Abidi, 2018) is distinctive from CGLLR algorithm in two aspects. First, the CGLLR algorithm, like any other conventional approaches to query rewriting (Pérez-Urbina, Motik, & Horrocks, 2009; Rosati & Almatelli, 2010), uses only the axioms of a $DL - Lite$ ontology (i.e., terminological component of the knowledge base) to reformulate a query. Then, the rewritten query will be evaluated over the extensional part of the ontology (i.e., assertional component). But as a data-driven approach, our PLausible Query Rewriting (PL-QR) algorithm considers both ontological constructs and assertional data to rewrite a query. Enrichment of the semantics of data with PL-OWL and merging data and domain ontology together in the form of a knowledge graph reinforce the advancement of OWL QL query rewriting over knowledge graphs using SPARQL 1.1.

Second, the rewriting rules in CGLLR algorithm are replaced with the plausible patterns as a set of inference constraints conducting the search strategy for the applicable semantics over the knowledge graph. In fact, the two algorithms are mainly different in their *principal search criteria* (step 6 of the plausible query rewriting algorithm, Algorithm 1).

The plausible query rewriting algorithm (Algorithm 1) starts with an initially failed query Q (Input 1), and ends with a set of conjunctive queries, R (Output). In addition to a failed query, the algorithm requires a set of preferred plausible patterns (Input 2) to limit the search for the applicable semantics to the relevant semantic data and terminological constructs (*TBox*) including both hierarchical and order-based relationships (Input 3).

We anticipate that the input ontology is in the $DL - Lite$ family since DLs of this family (i) are rich enough to express significant ontology languages, and (ii) query answering over $DL - Lite$ knowledge bases can be performed in a polynomial complexity. Even slight extensions (i.e., plausible OWL extension) to the logics of the family make query answering, worst case, *NLogSpace* in data complexity. Hence, the logics of the $DL - Lite$ family are effective description logic for implementing query answering over large data repositories (ABoxes) (Calvanese et al., 2007).

Algorithm 1- The proposed QR algorithm (Mohammadhassanzadeh, Raza Abidi, et al., 2017)

Input: (1) A query (triple pattern format),
(2) a set of plausible patterns $\pi \in \Pi$: $\{GEN, SPEC, SIM, DIS, FORT, INTP\}$,
(3) $DL - Lite$ TBOx $\mathcal{T}$ enriched with plausible OWL extension
Output: R, a set of rewriting queries.
1: R $= \{Q\}$;
2: **repeat**
3:   **foreach** $query$ Q $\in$ R **do**
4:     **foreach** $atom\ D\ in\ Q$ **do**
5:       **foreach** $pattern\ \pi \in \Pi$ **do**
6:         $\Delta = \phi(D, \pi, \mathcal{G})$
7:         **foreach** $D' \in \Delta$ **do**
8:           $Q' = \exists D'.Q(D \rightarrow D') \wedge \alpha(D, D')$;
9:           $R = R \cup \{Q'\}$;
10: **until** no unique query can be added to R;
11: **return** R;

Starting with the initial query, the algorithm (Algorithm 1) adds the query to the set R (step 1), which is empty when the algorithm starts. The algorithm repeats the following steps (steps 3 to 9) for each query in the set R until there is no new query to be added to R (step 10): for each atom[3] D in the query Q (step 4) and for each pattern π in the preferred plausible patterns (Input 2), the algorithm attempts to find (step 6) a set (Δ) of *applicable*

[3] Depending on the type of the query (i.e., yes/no question, factoid question), the user can identify the atom (i.e., subject or object of a triple statement) to be replaced.

plausible semantics (Definition 7) to the body atom via the *pattern matching function* ϕ. The functionality of the pattern matching function is explained in Algorithm 2.

Definition 7 (Applicable Plausible Semantic): for an atom D, atom D' is an applicable plausible semantic if (i) D' is semantically related to D ($\exists\alpha \in \mathcal{T}: \alpha(D, D')$), and (ii) α conforms with the plausible semantics (Σ) underlying the plausible patterns (Definition 2). For example, rdfs:subClassOf conducts generalization pattern, owl:instanceOf is applicable to specialization, owl:sameAs is used in similarity, and plowl:standsAfter is helpful to run a fortiori or interpolation.

For each *applicable plausible semantic* ($D' \in \Delta$), the algorithm constructs a new query Q' by replacing D with the atom D' (step 8) and will add the new query to the set R (step 9). The algorithm keeps rewriting new queries until there is no more unique query to be added (step 10).

Pattern matching function

The pattern matching function is basically a SPARQL query generator that queries a knowledge graph ($\boldsymbol{\mathcal{G}}$) to retrieve a set of applicable plausible semantics to atom $\boldsymbol{D}$ ($\boldsymbol{\Delta}$) conforming with the rationale of a plausible pattern ($\boldsymbol{\pi}$). Algorithm 2 demonstrates the steps of the pattern matching function.

The pattern matching function is implemented via a set of SPARQL queries that query the RDF repository based on (i) the body atom to be replaced, and (ii) the desired set of plausible patterns. Each *case* in Algorithm 2 demonstrates the construct of the SPARQL query corresponding to the relevant plausible pattern.

Algorithm 2. The pattern matching function

Input: (1) D: body atom of the query to change, (2) π: preferred plausible pattern, (3) $\mathcal{G}$: knowledge graph to explore

Output: Δ, a set of applicable semantics to atom D w.r.t π.

1: $\Delta = \{\emptyset\}$;
2: **switch** (π)
3: **case** GEN**:**
4: $\Delta = \phi(D, GEN, \mathcal{G}) = \{D' \mid [D\ \alpha\ D'][\alpha\ a\ \text{plowl: HierarchicalProperty}]\}$;
5: **break;**
6: **case** SPEC**:**
7: $\Delta = \phi(D, SPEC, \mathcal{G}) = \{D' \mid [D'\ \alpha\ D][\alpha\ a\ \text{plowl: HierarchicalProperty}]\}$;
8: **break;**
9: **case** SIM**:**
10: $\Delta = \phi(D, SIM, \mathcal{G}) = \{D' \mid [D\ \alpha\ D'][\alpha\ a\ plowl{:}HierarchicalSimProperty] \cup [D'\ \alpha\ D][\alpha\ a\ plowl{:}HierarchicalSimProperty]\}$;
11: **break;**
12: **case** DIS**:**
13: $\Delta = \phi(D, DIS, \mathcal{G}) = \{D' \mid [D\ \alpha\ D'][\alpha\ a\ plowl{:}\text{HierarchicalDissimProperty}] \cup [D'\ \alpha\ D][\alpha\ a\ plowl{:}\text{HierarchicalDissimProperty}]\}$;
14: **break;**
15: **case** AFORT**:**
16: $\Delta_{ML} = \phi(D, AFORTML, \mathcal{G}) = \{D' \mid [D'\ \alpha\ D][\alpha\ a\ plowl{:}\text{StandsBefore}] \cup [D\ \alpha\ D'][\alpha\ a\ plowl{:}\text{StandsAfter}]\}$;
17: $\Delta_{LM} = \phi(D, AFORTLM, \mathcal{G}) = \{D' \mid [D\ \alpha\ D'][\alpha\ a\ plowl{:}\text{StandsBefore}] \cup [D'\ \alpha\ D][\alpha\ a\ plowl{:}\text{StandsAfter}]\}$;
18: $\Delta = \Delta_{ML} \cup \Delta_{LM}$;
19: **break;**
20: **case** INTP**:**
21: $\Delta_A = \phi(D, INTPA, \mathcal{G}) = \{(D', D'') \mid [D'\alpha\ D][\ D\ \alpha\ D''][\alpha\ a\ plowl{:}\text{StandsBefore}]\}$;
22: $\Delta_D = \phi(D, INTPD, \mathcal{G}) = \{(D', D'') \mid [D''\alpha\ D][\ D\ \alpha\ D'][\alpha\ a\ plowl{:}\text{StandsAfter}]\}$;
23: $\Delta = \Delta_A \cup \Delta_D$;
24: **break;**
25: **return** Δ;

In the case of *generalization* (step 3 of Algorithm 2), the equivalent SPARQL query to the relation in step 4 will be as follows:

```
SELECT ?P ?O
WHERE
  {
    D ?P ?O.
    ?P rdf:type plowl:HierarchicalProperty
  }
```

Code 3- *Corresponding SPARQL query implementing generalization in the pattern matching function*

Code 3 searches for the applicable plausible semantic(s) to atom *D* (the body atom of the query that is going to be replaced) which are in a relationship with *D* via a *Hierarchical* property—e.g., *rdf:type, SemMedDB:ISA* (Code 2). Likewise, corresponding SPARQL queries to *specialization*, *similarity* and *dissimilarity* relationships in Algorithm 2 (steps 7, 10 and 13 respectively) are demonstrated in Table 8.

Table 8- Corresponding SPARQL queries implementing specialization, similarity and dissimilarity in the pattern matching function

Plausible Pattern	**Matching Function**	**Corresponding SPARQL query**

Specialization	$\phi(D, SPEC, \mathcal{G})$	`SELECT ?S ?P`
		`WHERE {`
		`    ?S ?P D.`
		`    ?P rdf:type plowl:HierarchicalProperty. }`

Similarity	$\phi(D, SIM, \mathcal{G})$	`SELECT ?P ?S ?O`
		`WHERE {`
		`    { D ?P ?O.`
		`      ?p rdf:type plowl:HierarchicalSimProperty. }`
		`    UNION`
		`    { ?S ?P D.`
		`      ?P rdf:type plowl:HierarchicalSimProperty. }`
		`  }`

Dissimilarity	$\phi(D, DIS, \mathcal{G})$	`SELECT ?P ?S ?O`
		`WHERE {`
		`    { D ?P ?O.`
		`      ?P rdf:type plowl:HierarchicalDissimProperty. }`
		`    UNION`
		`    { ?S ?P D.`
		`      ?P rdf:type plowl:HierarchicalDissimProperty. }`
		`  }`

In *similarity* and *dissimilarity* SPARQL queries, we are trying to consider both of the possible combinations of representing similar (or dissimilar) concepts – i.e., $sim(D, D')$ and $sim(D', D)$ – in order to have an exhaustive search of (dis)similar relationships. Hence, in the SPARQL queries the *union* of both possible combinations is required.

In the case of *a fortiori*, depending on the direction of an ordered relationship, a fortiori exploration on the graph can be conducted in two variants: *from more to less* or *from less to more*. Therefore, the corresponding SPARQL query in the pattern matching function combines the semantics resulting of both directions (Table 9). However, in the justification step of the plausibly inferred results, SeDan identifies the direction of the a fortiori pattern.

Table 9- Corresponding SPARQL query implementing two variants of a fortiori in the pattern matching function

A fortiori variation	Matching Function	Corresponding SPARQL query
More to Less	$\phi(D, AFORTML, \mathcal{G})$	`SELECT ?P ?S ?O WHERE { { ?S ?P D. ?P rdf:type plowl:StandsBefore. } UNION { D ?P ?O. ?P rdf:type plowl:StandsAfter. } }`
Less to More	$\phi(D, AFORTLM, \mathcal{G})$	`SELECT ?P ?S ?O WHERE { { D ?P ?O. ?P rdf:type plowl:StandsBefore. } UNION { ?S ?P D. ?P rdf:type plowl:StandsAfter. } }`

Like *similarity* and *dissimilarity* SPARQL queries, in each *a fortiori* variant (more to less and less to more) both of the possible combinations of representing an ordered-based relationship are queried and the results are combined via a *union*. Based on the nature of ordered associations, an ordered property can be represented in two directions. For example, the statement *"x is shorter than y"* can be represented as *shorter(x,y)*, while the triple *taller(y,x)* implies the same impression. Therefore, using the axioms introduced in the plausible OWL extension, there are two alternatives to represent an ordered relationship: $StandsBefore(x, y) \equiv StandsAfter(y, x)$. Similarly, the ordered relationships between three concepts can be represented via either $StandsBefore$ (ascending order) or $StandsAfter$ (descending order) axioms:

$$StandsBefore(x, y), StandsBefore(y, z) \equiv StandsAfter(z, y), StandsAfter(y, x)$$

Hence, in the SPARQL query for *interpolation*, both of the possible alternatives will be queried, and the results will be combined to further be used in the rewriting algorithm later (Table 10).

Table 10- Corresponding SPARQL query implementing two variants of interpolation in the pattern matching function (LB: lower bound, UB: upper bound)

Interpolation variation	Pattern Matching Function	Corresponding SPARQL query

Ascending order	$\phi(D, INTPA, \mathcal{G})$	`SELECT ?P ?LB ?UB`
		`WHERE {`
		`    ?LB ?P D.`
		`    D   ?P ?UB.`
		`    ?P rdf:type plowl:StandsBefore. }`

Descending order	$\phi(D, INTPD, \mathcal{G})$	`SELECT ?P ?LB ?UB`
		`WHERE {`
		`    ?UB ?P D.`
		`    D   ?p ?lb.`
		`    ?P rdf:type plowl:StandsAfter. }`

4.4 SeDan: a semantics-based data analytics framework

The sections above, PLausible extension to OWL (PL-OWL) and PLausible Query Rewriting algorithm (PL-QR), introduced the solutions that, along with the enrichment of underlying semantics with PL-OWL, can address the challenges of implementing plausible reasoning over knowledge graphs. However, to establish the *act of plausible reasoning* (Definition 6) and accomplish semantics-based data analytics in real settings, it is required to integrate the solutions together in one framework manifesting a purposeful plausible reasoner. In this regard, we developed a SEmantic-based Data ANalytics (SeDan) framework that implements a plausible reasoner to infer new knowledge from RDF knowledge graphs.

Figure 3 demonstrates the reasoning approach of the SeDan framework when a new query arrives. In the first step (step 1), the original query, with no (plausible) manipulation, will be asked from the knowledge graph to provide the deductive answers existing in the data (step 2). If an answer(s) is retrieved, it will be reported as a deductive solution(s) to the query (step 2.1). But, in the case of a failed query (step 2.2), the plausible reasoner will be invoked (i.e., in line with what Definition 6 suggests).

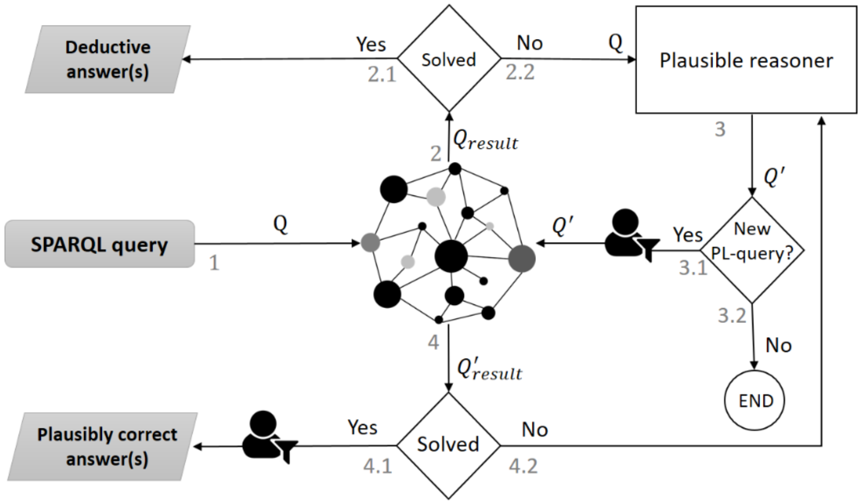

Figure 3- Reasoning mechanism of SeDan framework (the numbers show the sequence of the steps)

The plausible reasoner retrieves the applicable semantics from the knowledge graph to reformulate (step 3) the given query into a plausible query(ies). The subsequent plausible queries (if there is any), along with the supporting semantics that lead to the resulting plausible queries (i.e., the plausible path), will be provided to the *user* (step 3.1). Then, the user investigates the queries, validates the rationality (meaningfulness) of the plausible queries, and filters out the meaningless queries. It is worth to mention that at this point, SeDan is considered as a proof of concept (i.e., a research tool and not a diagnostic or decision support tool) that answers queries in research environments. Hence, the *users* of the system are researchers and scientists who possess the domain knowledge and can verify the meaningfulness of the plausible paths and evaluate the quality of the plausible queries generated by the PL-QR algorithm.

Afterwards, the *good* (i.e., acceptable) plausible queries (if there is any) will be asked from the knowledge graph with the aim of providing a plausible answer(s) to the query. If there is any statement in the knowledge graph satisfying the graph pattern of the plausible query, then it will be presented as a plausible answer(s) (step 4.1); otherwise, the failed query will be considered for further rewriting by the plausible reasoner (step 4.2). The rewriting of the queries continues until there is no new plausible query to be generated (i.e., or the reasoning process is halted by the user). Figure 4 further elaborates on the sequence of the steps in the process of plausible reasoning. The domain expert is engaged in the process in two steps: (i) the evaluation of the acceptability of plausible queries, and (ii) the validation of plausibly inferred answer(s).

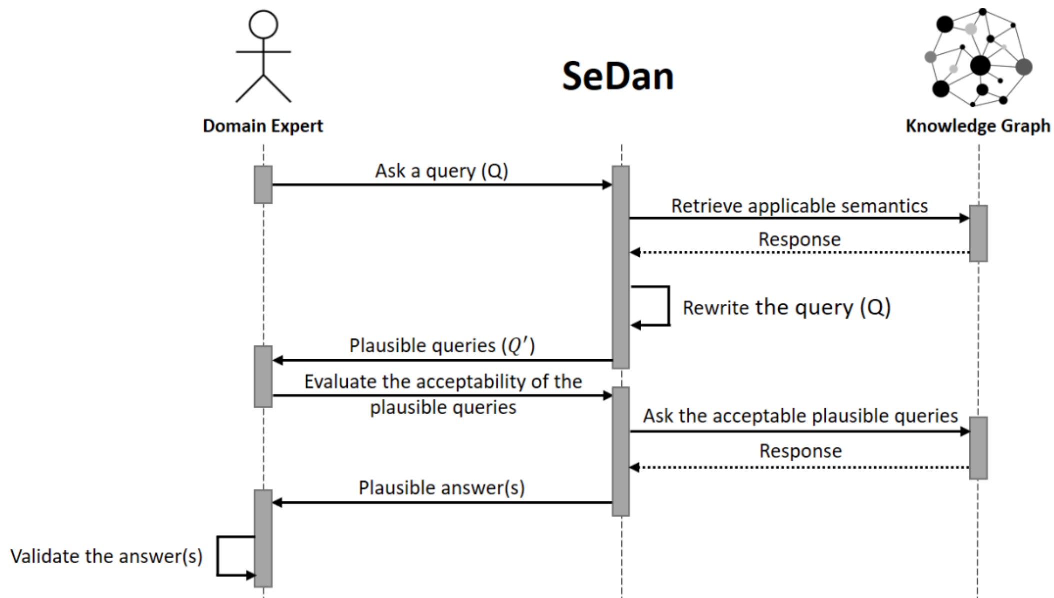

Figure 4- Plausible reasoning sequence diagram

The domain expert needs to evaluate the acceptability of the queries and the validity of the inferred plausible answers due to the uncertain nature of plausible reasoning. As it was discussed in Chapter 3, a plausible argument is subjective. In addition, the plausible reasoner generates all the possible plausible queries, which not all of them are clinically correct or relevant to the original query. Hence, a plausible answer and its justification should be in line with the purpose of the question and rationale to the user who asked the question (Figure 4).

Table 11- Examples of generated plausible queries for different questions, which their acceptability was evaluated by the domain expert (for each question there are more plausible queries, but only few of them are provided here)

Question	Plausibly generated queries	Patterns involved	Plausible Acceptable
Is Herceptin of potential use in the treatment of prostate cancer?	`ASK { Antibodies TREATS Prostate_cancer }`	{GEN}	✓
	`ASK { Therapeutic_agent TREATS Prostate_cancer }`	{GEN}	✓
	`ASK { agonists TREATS Prostate_cancer }`	{SIM-OR}	O
	`ASK { Zoladex TREATS Prostate_cancer }`	{GEN, SPEC}	O
	`ASK { Radioimmunoconjugates TREATS Prostate_cancer }`	{GEN, SPEC}	O
	`ASK { inhibitors TREATS Prostate_cancer. Antibodies TREATS Prostate_cancer.}`	{GEN, INTP}	O
	`ASK { inhibitors TREATS Prostate_cancer . receptor TREATS Prostate_cancer. }`	{GEN, INTP}	O
What is the treatment of acute myocarditis?	`SELECT ?x WHERE { ?x TREATS Breakthrough_pain. }`	{GEN, SIM-HR, GEN, SPEC}	✓
	`SELECT ?x WHERE { ?x TREATS Ventricular_Dysfunction,_Left. }`	{GEN, SIM-HR, AFORT-ML, INTP}	O
	`SELECT ?x WHERE { ?x TREATS Infarction.}`	{GEN, SIM-HR, AFORT-ML, INTP}	O
What organism causes Woolsorter's disease?	`SELECT ?x WHERE { ?x semp:CAUSES semr:Rupture,_Spontaneous. }`	{SIM-HR, SIM-HR, GEN, AFORT-ML}	O
	`SELECT ?x WHERE { ?x semp:CAUSES semr:Brain_Edema. }`	{SIM-HR, SIM-HR, GEN, SIM-OR}	✓
	`SELECT ?x WHERE { ?x semp:CAUSES semr:Necrosis. }`	{SIM-HR, SIM-HR, GEN, SIM-OR}	O
Matuzumab has been tested for treatment of which cancers?	`SELECT ?x WHERE { Cisplatin TREATS ?x; }`	{GEN, SIM-HR, SIM-OR}	O
	`SELECT ?x WHERE { Pharmaceutical_Preparations TREATS ?x. }`	{GEN, SIM-HR, GEN}	✓
	`SELECT ?x WHERE { Aspartate_Transaminase TREATS ?x. }`	{GEN, SIM-HR, SIM-OR}	✓
	`SELECT ?x WHERE { Gene_Transduction_AgentTREATS ?x. }`	{GEN, SIM-HR, SIM-OR}	✓

Table 11 presents the plausible queries (few plausible queries out of many) generated for four medical questions that the Semantic MEDLINE database cannot answer them deductively. For example, in the case of the question asking if *"Herceptin is of potential use in the treatment of prostate cancer?"*, out of 7 plausible queries (the table does not include all the generated plausible queries) only 2 of them are clinically acceptable. Even in the case of an acceptable plausible query, not all the plausibly inferred answers are relevant and correct. For example, *Lidocaine*, *Ibuprofen* and *Diclofenac* (Table 12) are answers to an acceptable plausible query (*?x TREATS Breakthrough_pain*), but only *Ibuprofen* was found plausibly correct, an the other two are found incorrect/irrelevant with regard to the concept in the original question, *acute myocarditis*. Chapter 5 will elaborate on the evaluation process in detail and discusses further why a plausible query or a plausible answer is found unacceptable or irrelevant.

Table 12- Examples of plausible answers for acceptable plausible queries, which their correctness is evaluated by the domain expert

Plausible query	Plausible answers	Plausible acceptable
`SELECT ?x WHERE { ?x TREATS Breakthrough_pain. }`	`Lidocaine`	O
	`Ibuprofen`	✓
	`Diclofenac`	O
`SELECT ?x WHERE { ?x semp:CAUSES semr:Brain_Edema. }`	`Pertussis_Vaccine`	O
	`Bacillus_infection`	✓
	`Agent`	O
	`Liver_Failure`	O
`SELECT ?x WHERE { Aspartate_Transaminase TREATS ?x. }`	`Gastric cancer`	O
	`Diabetes`	✓
	`Syphilis`	✓
	`Hepatitis B`	✓

4.4.1 SeDan's architecture

The plausible reasoner, the core component of the SeDan framework, develops the plausible patterns by manipulating the underlying graph directly with SPARQL query rewriting using OWL 2 QL and the introduced PL-OWL constructs. As Figure 5 demonstrates (and discussed before), the plausible reasoner is comprised of two modules: the pattern matching function and the query rewriting algorithm.

The pattern matching function exploits the built-in constructs of OWL QL and the constructs of PL-OWL to develop the rationale behind the plausible patterns when retrieving the applicable semantics. Then, using the retrieved semantics, the query rewriting algorithm transforms a given query to a plausible version. In SeDan (Figure 5), the knowledge graph is the source for both answering and rewriting the queries. Conforming with the notion, the knowledge graph in SeDan combines both the domain knowledge (i.e., expressed in the form of (DL) ontologies) and the domain data (i.e., represented in RDF). Thus, the knowledge graph includes the material (i.e., assertional data) to answer a SPARQL query, while it provides the required semantics (i.e., retrieved from ontological constructs or semantic data) to conduct the query rewriting.

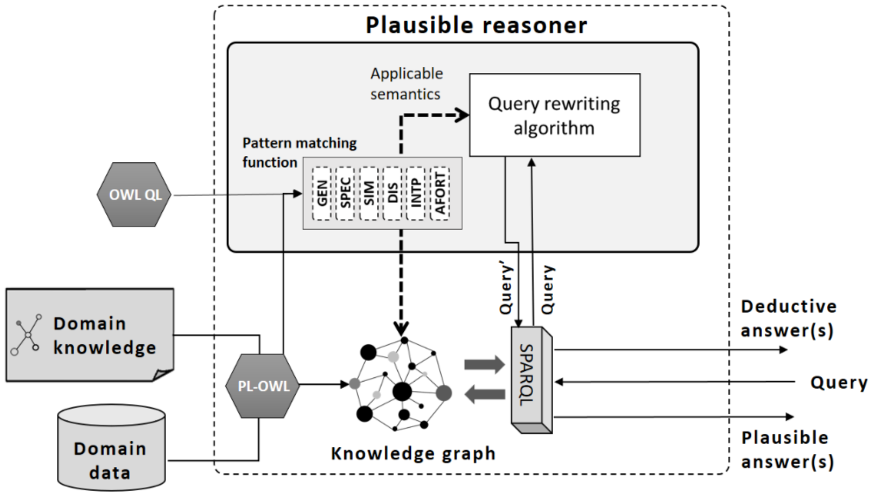

Figure 5- The architecture of the SeDan framework

The following section provides some case studies elaborating on the functionality of our plausible reasoner in the SeDan framework.

4.5 Case studies

To demonstrate the functionality of SeDan's query rewriting algorithm, we provide two case studies where we attempt to answer two questions from BioASQ challenges[4] (Tsatsaronis, Balikas, Malakasiotis, & et. al., 2015) using DrugBank[5] (Knox et al., 2011; Law, Knox, Djoumbou, & Jewison, 2014; Wishart et al., 2018), Disease Ontology[6] (Kibbe, Arze, Felix, & Mitraka, 2015; Schriml et al., 2012) and Semantic MEDLINE[7] database[8]

[4] http://www.bioasq.org/
[5] DrugBank datasets are released under a Creative Common's Attribution-NonCommercial 4.0 International Public License
[6] Disease ontology files are available on http://disease-ontology.org/ under the Creative Commons license.
[7] The National Library of Medicine freely provides PubMed/Medline Data (more information: NLM Copyright Information).
[8] Available at https://skr3.nlm.nih.gov/SemMedDB/index.html

(Kilicoglu, Shin, Fiszman, Rosemblat, & Rindflesch, 2012; Rindflesch, Kilicoglu, & Fiszman, 2011). These materials and resources will be explained further in the Evaluation section.

4.5.1 Example 1: statins cause diabetes?

In this example, we provide a case study to elaborate on how SeDan, and in particular, the plausible reasoner leverage assertional data and ontological constructs to transform a (deductively) failed query to a new plausible query. For a question asking *"Do statins cause diabetes*?" (BioASQ challenge, Task 2b), the initial SPARQL syntax of the question can be written as below:

Initial SPARQL query: @PREFIX sem: <https://skr3.nlm.nih.gov/SemMed#> ASK { "statins" sem:causes "diabetes" }	Answer: No

Code 4- *Initial query answering if statins cause diabetes*

Conventional deductive reasoning over the SemMedDB returns *'No'*, since there is no matching triple that unifies the query. However, by leveraging the semantics of Drug Bank and SemMedDB, we know:

```
("Pravastatin", db:substructure, "statins")        (1)
("Pancreatitis", sem:precedes, "Diabetes")         (2)
("Diabetes", sem:precedes, "Hyperglycemia")        (3)
```

Figure 6- Relevant semantics to the concepts in the query – substructure represents hierarchical relationships in DrugBank and precedes characterizes order-based relationships in SemMedDB

In the statements above, the *sub-structure* predicate (Figure 6, semantic 1) represents a hierarchical relationship (Code 2) implying *Pravastatin* (DB00175) belongs to a class of medications known as *statins*. The *precedes* predicate, as an *plowl:OrderedProperty*, shows a "*sequence" (order)* of phenomena in the SemMedDB (Code 2). Semantics 2 and 3 in Figure 6 represent the order of three diseases, *Pancreatitis* → *Diabetes* → *Hyperglycemia*, that can happen sequentially. Using the semantics above, the query rewriting algorithm leverages the specialization and interpolation patterns to transform the initial query to the plausible query below:

Rewritten SPARQL query:	Plausible Answer:

`PREFIX db: <https://www.drugbank.ca/drugs#>`
`PREFIX sem: <https://skr3.nlm.nih.gov/SemMed#>`
`ASK`
`{ "Pravastatin" sem:causes "Pancreatitis".`
`  "Pravastatin", sem:causes, "Hyperglycemia". }`
	(Yes, {SPEC, INTP})

Code 5- Rewritten query answering if statins cause diabetes

By posing the new query over SemMedDB (i.e., considering that the expert found the plausible query as a good (meaningful) plausible query), we will get a plausible positive answer of *Yes*, which is inferred via a combination of *specialization* and *interpolation* patterns. *Pravastatin,* as a type of *statins* (a hierarchical relationship) could cause *Pancreatitis* and *Hyperglycemia*, which are two diseases that *occur* before and after *diabetes*, respectively (an ordered relationship). Based on the rationale behind the *specialization* pattern that *"when something is true about a class/entity, it might be true about its sub-classes as well"* and the logic of *interpolation pattern* that *"if something is true about two (stages of) phenomena, then it might be true for any phenomenon in between",* we could say *statins* plausibly *causes diabetes*, since one of its instances, *Pravastatin, causes* diseases that are prior and subsequent to *diabetes*.

It should be noted that the example above represents only one plausibly rewritten query out of many possible queries. Depending on the existing axioms in the ontology and supporting semantics relevant to the concept(s) in the query, the PL-QR algorithm continues rewriting new queries until there is no more unique query to be retrieved.

4.5.2 Example 2: Herceptin treats Prostate Cancer?

In this case study, we are asking another Yes/No question, *"Is Herceptin of potential use in the treatment of prostate cancer?"* (BioASQ challenge, Task 2b), from the SemMedDB. Making use of the existing triples in the database, there is no matching triple unifying the question. Consequently, the answer will be 'No'. The initial SPARQL syntax of the question is as bellow:

Initial SPARQL query:	Answer:
@PREFIX sem: <https://skr3.nlm.nih.gov/SemMed#> ASK { "Herceptin" sem:treats "Prostate cancer" }	No

Code 6. Initial query answering if Migalastat treats Fabry Disease

Utilizing the Disease ontology axioms (DOID:10286) and the existing triples in SemMedDB, we know:

```
("Herceptin", sem:treats, "Malignant neoplasms")                (1)
("Malignant neoplasms", sem:occurs_in, "Prostate carcinoma")  (2)
("Prostate carcinoma", do:isa, "Prostate cancer")              (3)
```

Figure 7- Rewritten triple using the ontology axiom

In the triples above, the *treats* predicate (Figure 7, semantic 1) shows a disease, *malignant neoplasms*, that could be treated by *Herceptin*. The *occurs_in* relationship (Figure 7, semantic 2) characterizes the *"order"* of the occurrence of two phases of a disease: *malignant neoplasms* and *prostate carcinoma*. The *is_a* relationship (Figure 7, semantic 3) represents a hierarchical relationship between two diseases, *prostate carcinoma* and *prostate cancer*. Using the semantics above, the PL-QR algorithm exploits the specialization and a fortiori patterns, to transform the initial query to the expanded query below:

Rewritten SPARQL query:	Plausible Answer:
PREFIX do: < http://disease-ontology.org/term#> PREFIX sem: <https://skr3.nlm.nih.gov/SemMed#> ASK { "Herceptin" sem:treats " Malignant neoplasms". "Malignant neoplasms", sem:occurs_in, "Prostate carcinoma". "Prostate carcinoma", do:isa, "Prostate cancer"}	(Yes, {SPEC, AFORT})

***Code 7-** Rewritten query answering if Migalastat treats Fabry Disease*

After the approval of the query by the expert as an acceptable plausible query and asking the new query from SemMedDB, we will get a plausible positive answer that is inferred via the specialization and a fortiori patterns. The inference above means *Herceptin* could *treat prostate cancer*, as *Herceptin* could *treat malignant neoplasms* that is an *earlier* phase (ordered relationship) of *prostate carcinoma*, which *is a type of* (hierarchical relationship) *prostate cancer*. In other words, *Herceptin* could *plausibly treat prostate cancer* as it is administered to some prior phases of the disease.

4.6 Summary

This chapter started with providing a formal description of plausible reasoning addressing one of the key research challenges of the work. The formal definition of plausible reasoning and its components suggest an unambiguous understanding of the notion that (i) maps with the definition of plausible reasoning in the literature and the theory and (ii) is applicable to computer systems and conforms to the Semantic Web inspired approach that we have taken (i.e., especially with the graph representation of knowledge).

The formal definitions introduced in this chapter clarify the requirements and challenges of implementing plausible reasoning over knowledge graphs. The expressivity of the knowledge representation formalism and the flexibility of the reasoning paradigm are two issues that challenge the development of the plausible patterns and the representation of the semantics that they are applied to. This chapter showed our solutions, including plausible extension to OWL (PL-OWL) and plausible query rewriting algorithm (PL-QR), can address these issues.

PL-OWL introduces new constructs to existing OWL axioms enabling it to capture ordered relationships. PL-QR captures and develop the rationale of plausible patterns in the form of inference constraints conducting the query rewriting. We introduced the SeDan framework that integrates the solutions in one solid framework manifesting the concept of semantics-based data analytics.

Chapter 5: Evaluation and Experimental Results

An effective evaluation allows for understanding of the system's usefulness, precision, feasibility, as well as the potential avenues to improve (Milstein & Wetterhall, 1999). In the medical world, doctors are required to answer intelligent and complex questions based on large amounts of health data, which is usually sparse, noisy, incomplete and uncertain.

To evaluate the efficacy of SeDan, we aim to tailor a practical medical setting, tied with the routine procedure of answering medical questions. Inspired by the Lehigh University Benchmark (LUBM) (Guo, Pan, & Heflin, 2005), SeDan will be evaluated in three dimensions: *functionality* of the plausible reasoner, *correctness* of the plausible answers, and *cost-effectiveness* of the system.

Within the functionality evaluation (Figure 8), we try to assess if the plausible reasoner was implemented properly according to its primary objectives and operates as it was intended to. The functionality evaluation focuses on the navigation task of the reasoning engine over the knowledge graph, without considering the acceptability or adequacy of the answers. We expect plausible patterns, alone or in combination with other patterns, explore the knowledge graph, provide plausible answer(s) and extend the query answering coverage of the knowledge graph.

The second aspect of the evaluation paradigm, correctness of the plausible answers, investigates the acceptability of the plausibly inferred answer(s) and the validity of the reasoning processes reaching the answer(s).

Figure 8- SeDan's evaluation framework

Definition 8 (Plausibly correct answer): A plausibly inferred answer is deemed as *plausibly correct* if it is (i) reached via a meaningful, relevant set of semantics (*acceptable plausible query*) and (ii) validated by a medical domain expert. Plausibly correct answers are not evaluated against any evidence-based resources and are not ranked.

Definition 9 (Acceptable plausible query): An *acceptable plausible query* is a query that is generated by the plausible query rewriting algorithm via a meaningful set of semantics (aka. *plausible path*) that are justifiable from clinical standpoint.

Hence, the correctness evaluation of the answers is a two-step process: (i) a domain expert will evaluate the plausible queries and filter out the meaningless queries, (ii) acceptable plausible queries will be posed to the knowledge graphs and the subsequent answers will be evaluated by the domain expert.

Moreover, the cost-effectiveness evaluation aims to assess efficient use of time and resources (Jafarpour, Raza Abidi, Van Woensel, & Raza Abidi, 2019) and determines the practicality of SeDan in real world settings. The *performance* and *feasibility* of SeDan in plausible exploration of the knowledge graph (i.e., particularly the query rewriting algorithm) and discovery of the correct answers are studied via three measures: *retrieval time of applicable plausible semantics*, *rewriting time of a plausible query*, *time to reach the first approved answer* (also know as *response time*).

This chapter introduces the essential elements of SeDan's evaluation, presents the evaluation design, how it is conceived and conducted. Furthermore, the chapter will elaborate on the required materials to performing the evaluation. Ultimately, this chapter presents the experimental results and discusses the criteria being measured.

5.1 Experiment design

Our approach to perform the experiments includes two main phases. The first phase starts with posing the questions to the knowledge graph (RDF repository) in the form of plain SPARQL queries, without plausible reasoning. The outcome of this phase will divide the questions to two sets: initially answered questions and unanswered questions.

The concept(s) in the initially answered questions are included in the knowledge graph and their correct (gold standard) answers (i.e., yes/no or facts) are successfully retrieved via SPARQL queries. While, the unanswered ones are those questions that the system either does not return any responses (i.e., because the knowledge graph does not include the necessary concepts or the answer) or provides wrong answers (i.e., compared to the gold standard answers or verified by an expert).

In the second phase, the unanswered questions were posed again to SeDan, but with the plausible reasoning activated. The plausible reasoning engine receives the query and initializes the iterative process of the query rewriting algorithm, leveraging the existing semantics in the knowledge graph (i.e., including domain ontologies and data). The resulting plausible queries (if any) were provided to the domain expert (e.g., health care practitioners). The plausible queries that are deemed as acceptable (i.e., based on the plausible path) were posed to the knowledge graph (i.e., RDF triple store) with the aim of retrieving plausible answers. Figure 9 depicts the flow and design of the evaluation practice and the modules involved at each phase.

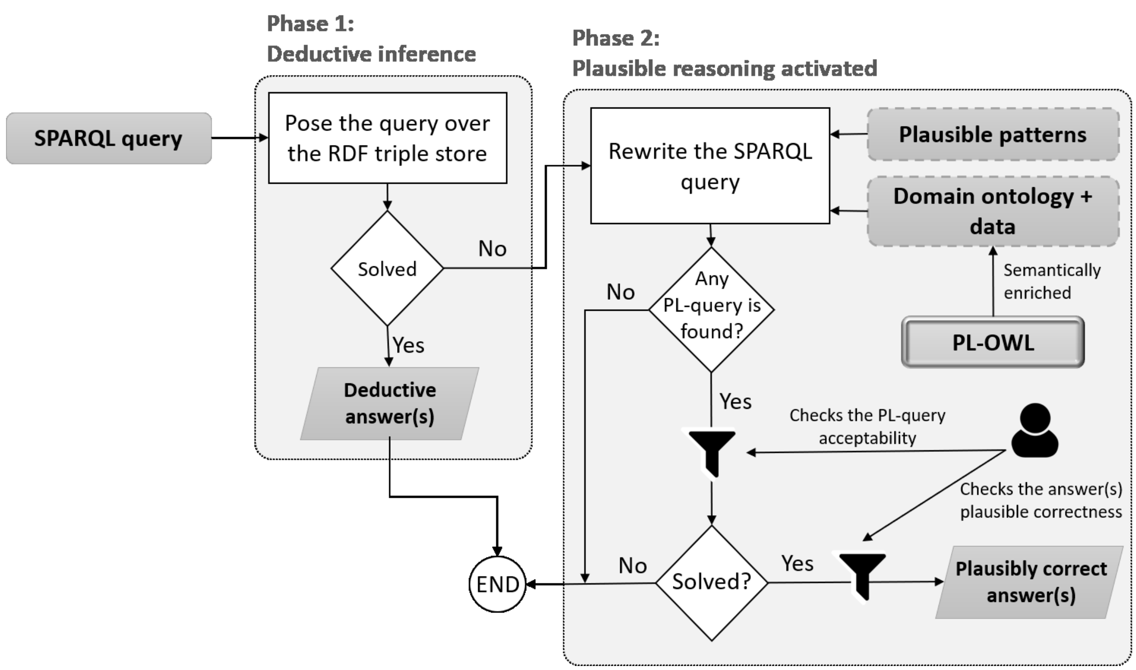

Figure 9- The workflow and design of the experiment

5.1.1 Experiment environment

The experiments were performed on a desktop computer with the system configuration (both hardware and software) as follows:

Hardware

- Operating System: Windows 10 Home (64-bit)
- CPU: Intel® Core™ i7-4770 CPU @ 3.40GHz
- RAM: 12.0 GB Dual-Channel DDR3
- HDD: Seagate 2 TB SATA-III 6.0Gb/s (7200 RPM)

- Graphics: 1023 MB NVIDIA GeForce GTX 645 (NVIDIA)

Software

- Java JDK 1.7.0
- Java JRE 1.8.0
- Eclipse - Standard Luna-R (win32-×86)
- GraphDB

5.2 Resources and Materials for the Experiment

An effective implementation of the designed experiment requires a specific set of resources, including: (i) a set of medical questions; (ii) a large health data set; and (iii) background domain ontologies, complementary to the health data.

The medical questions that are going to be answered should be challenging enough that require certain knowledge or skills to answer, yet resolvable leveraging latest medical knowledge available—i.e., to be able to compare the plausibly inferred answers with medically approved responses. Additionally, it is desirable that the experimental questions are convertible to SPARQL queries with no complications, since the complexity of parsing natural language questions to SPARQL queries is not the focus of this study. In addition, to avoid any bias in the evaluation of the system, the questions should be provided from external sources and the domain expert, who is involved in the development and evaluation of the system, should not be engaged in the process of selecting the questions.

Pertinent with the experimental questions, we need to load the RDF repository with a large health *data* set, which (i) can provide answers to the medical queries, (ii) is representable as a knowledge graph, and (iii) is preferably incomplete—i.e., here, the term *data* indicates assertions about instances, like ABox in Description Logic (DL) terminology.

To support the reasoning engine with the required semantics to conduct plausible reasoning, SeDan could be enriched with the background domain knowledge (like TBox in DL terminology) relevant to the experimental questions and the underlying data. The domain knowledge could be a set of ontologies that incorporate the concepts in questions and their associations; semantically richer ontologies would enhance the efficiency and performance of the plausible reasoning. In our experiment, we asked questions from BioASQ challenges from Semantic MEDLINE database (Rindflesch et al., 2011), DrugBank (Law et al., 2014) and Disease Ontology (Kibbe et al., 2015), utilizing the query rewriting algorithm to rewrite the queries with no initial answers. The sources, the data processing, and the selection of medical questions are discussed below (Figure 10).

BioASQ Medical Questions. *BioASQ* challenges (Tsatsaronis et al., 2015) are a series of competitions (2013-2017) on large-scale biomedical semantic indexing and question answering. The purpose of the challenges is to assess the capability of machines to semantically index very large numbers of healthcare and life science publications and ontologies to compose brief and easy to understand answers to real-life biomedical questions.

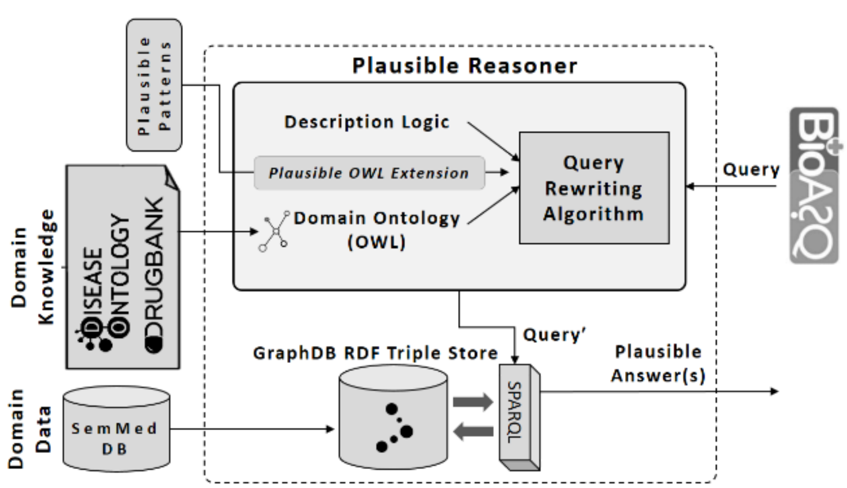

Figure 10- SeDan framework and the design of the experiment

BioASQ questions are formulated by European biomedical experts, reflecting a variety of real-life inquiries. The questions belong to 4 distinct categories: yes or no, factoid, list, and summary questions. While the BioASQ questions belong to a variety of contexts, biology, pharmacology, etc., for our evaluation we focused on those BioASQ questions that ask (i) for *yes/no* or *factoid* answer, and (ii) about *treatment* or *diagnoses*. We did not work with summary questions as they require Natural Language Processing techniques to prepare the answer and that is not within the scope of this study. Also, some of the questions include qualitative terms (i.e., the most known bacterium) that are not easily transformed into SPARQL queries.

To ask the BioASQ questions from SeDan we were required to translate the questions into SPARQL queries. In this regard, first, the medical entities and semantic relationships relevant to the questions were retrieved from the knowledge base. Then, the type of the expected answer (*yes/no* or *factoid*) was identified. Having these essentials ready, the equivalent SPARQL query (*ask* or *select*) would be constructed by manifesting the header and the body of the query respectively. The details of the retrieved questions from BioASQ Task 5, the most recent challenge of the series, is further elaborated in the section **5.3 Statistics on the experimental questions**.

Semantic MEDLINE Database. *Semantic MEDLINE database* (SemMedDB) (Kilicoglu et al., 2012; Rindflesch et al., 2011) is a significant endeavor to facilitate healthcare and life science studies by providing a comprehensive resource of structured semantic predications. The version of SemMedDB that was deployed in the experiment contains over 89 million records (as subject-predicate-object triples) extracted from over 26 million biomedical publications (as of Apr. 30, 2016).

Table 13- National Library of Medicine (NLM) semantic groups

Semantic Group	Examples of Semantic Types in each group
Activities and Behaviors	Daily or Recreational Activity; Event; Governmental or Regularity Activity; Machine Activity; Occupational Activity; Social Behavior
Anatomy	Body Location or Region; Organ; Body Space or Junction; Body Substance; Cell; Cell Component; Embryonic Structure; Tissue
Chemical and Drugs	Amino Acid, Peptide, or Protein; Antibiotic; Clinical Drug; Eicosanoid; Enzyme; Lipid; Nucleic Acid, Pharmacologic Substance; Steroid; Vitamin
Concepts and Ideas	Conceptual Entity; Functional Concept; Intellectual Product; Regulation or Law; Spatial Concept; Temporal Concept
Devices	Drug Delivery Device; Medical Device; Research Device
Disorders	Acquired Abnormality; Anatomical Abnormality; Cell or Molecular Dysfunction; Disease or Syndrome; Mental or Behavioral Dysfunction
Genes and Molecular Sequences	Amino Acid Sequence; Carbohydrate Sequence; Gene or Genome; Molecular Sequence; Nucleotide Sequence
Geographic Areas	Geographic Area
Living Beings	Age Group; Alga; Fish; Fungus; Human; Mammal; Plant; Population Group; Professional or Occupational Group; Reptile; Vertebrate; Virus
Objects	Entity; Good; Manufactured Object
Occupations	Biomedical Occupation or Discipline; Occupation or Discipline
Organizations	Health Care Related Organization; Organization; Professional Society; Self-help or Relief Organization
Phenomena	Biologic Function; Human Caused Phenomenon or Process; Laboratory or Test Result; Natural Phenomenon or Process; Phenomenon or Process
Physiology	Cell Function; Clinical Attribute; Genetic Function; Mental Process; Molecular Function; Organ or Tissue Function; Organism Function;
Procedures	Diagnostic Procedures; Educational Activity; Health Care Activity; Laboratory Procedure; Therapeutic or Preventive Procedure

In SemMedDB, the concepts (subjects and objects of the predications) belong to about 120 Unified Medical Language System (UMLS) semantic types (i.e., activity, vitamin, etc.) that are grouped into 15 National Library of Medicine (NLM) semantic groups (Table 13). Predicates are distributed among 34 relationships (i.e., causes, occurs in, etc.) and an additional 27 negation relationships (Table 14). In our experiment, because of the nature of the questions retrieved from BioASQ challenges, we use only 3 semantic groups: disorders (DISO), chemicals & drugs (CHEM), and genes & molecular sequences (GENE). Hence, any combination of these semantic groups, including 6 types of predications; DISO-DISO, DISO-CHEM, DISO-GENE, CHEM-CHEM, CHEM-GENE and GENE-GENE, were extracted (Tao, Zhang, Jiang, Bouamrane, & Chute, 2012). The resulting RDF repository contains over 11 million semantic predications.

Table 14- List of the predicates from Semantic Medline DB

Predicate	Description
Administered to	Given to patient, when no assertion is made that the substance is being given as treatment.
Affects	Produces a direct effect on. Implied here is the altering or influencing of an existing condition, state, situation, or entity.
Associated With	Has a relationship to (gene-disease relation)
Augments	Expands or stimulates a process
Causes	Brings about a condition or an effect. Implied here is that an agent, such as for example, a pharmacologic substance or an organism, has brought about the effect.
Coexist with	Occurs together with, or jointly.
Compared With	-
Complicates	Causes to become more severe or complex, or results in adverse effects
Converts to	Changes from one form to another (both substances)
Diagnoses	Distinguishes or identifies the nature or characteristics of
Different from	-
Different than	-
Disrupts	Alters or influences an already existing condition, state, or situation. Produces a negative effect on.
Higher from	
Higher than	
Inhibits	Decreases, limits, or blocks the action or function of (substance interaction)
Interacts with	Substance interaction
IS-A	The basic hierarchical link in the UMLS Semantic Network. If one item “is a” another item, then the first item is more specific in meaning than the second item.
Location of	The position, site, or region of an entity or the site of a process.
Lower than	
Manifestation of	That part of a phenomenon which is directly observable or concretely or visibly expressed, or which gives evidence to the underlying process.
Method of	The manner and sequence of events in performing an act or procedure.
Occurs in	Has incidence in a group or population.
Part of	Composes, with one or more other physical units, some larger whole. Includes component of, division of, portion of, fragment of, section of, and layer of.
Precedes	Occurs earlier in time. This includes antedates, comes before, is in advance of, predates, and is prior to.
Predisposes	To be a risk to a disorder, pathology, or condition.
Prevents	Stops, hinders or eliminates an action or condition.
Process of	Disorder occurs in (higher) organism.
Produces	Brings forth, generates or creates. This includes yields, secretes, emits, biosynthesizes, generates, releases, discharges, and creates.
Same as	-
Stimulate	Increases or facilitates the action or function of (substance interaction).
Than as	-
Treats	Applies a remedy with the object of effecting a cure or managing a condition.
Uses	Employs in the carrying out of some activity. This includes applies, utilizes, employs, and avails.

DrugBank. *DrugBank* (Kibbe et al., 2015; Law et al., 2014; Schriml et al., 2012) is a comprehensive database of biochemical and pharmacological information about drugs and drug targets. Each drug entry includes extensive information on properties, structure, and biology of the drugs. In the current setting of SeDan, we are exploiting *DrugBank* version 4.5.0, of which the RDF representation contains over 3.8 million predications in total.

Disease Ontology. The *Human Disease Ontology* (DO) (Kibbe et al., 2015; Schriml et al., 2012) is a standardized ontology for both common and rare human diseases. The *Disease Ontology* semantically integrates disease and medical vocabularies across disparate biomedical resources; MeSH, ICD, NCI's thesaurus, SNOMED and OMIM. The most up to date version of the DO contains 203,125 semantic predications.

GraphDB. The aforesaid materials and sources are stored to and queried via *GraphDB*[9] RDF triple store (http://graphdb.ontotext.com/). GraphDB is a graph database with RDF and SPARQL support. Its capabilities for semantic inferencing, efficient handling of massive volumes of data, real-time inferencing and support of quadruples make *GraphDB* an appropriate tool for an SPARQL endpoint in SeDan architecture.

5.3 Statistics on the experimental questions

As explained in the previous chapter, in the experiments, we focused on the questions of BioASQ challenges that (i) are confined within the domains of *treatment* or *diagnosis* (i.e., common medical questions that doctors are confronted with); and (ii) ask for *yes/no* or *factoid* answers (i.e., types of questions that SeDan answers). Based on two initial criteria for selecting the questions, 114 questions[10] were retrieved: 61 questions asking about *causes* of diseases and 53 questions asking about *treatments*. 83 questions look for *factoid* answers and the remaining 31 questions expect *yes* or *no* as the answer.

Table 15 presents the distribution of the retrieved questions.

---

[9] GraphDB Free is available under an RDBMS-like free license. It is free to use but not open-source (more information here).
[10] The list of questions is available at https://tinyurl.com/y7fr3yvd

Table 15- Details of the retrieved questions from BioASQ series, Task 6-Phase B question sets

Source	Release date	No. of questions	Retrieved questions				
			Total	Question domain		Questions type	
				Causes	Treatments	Factoid	Yes/No
Train	5-Dec-17	2251	69	36	33	58	11
Testset1	8-Mar-18	100	9	3	6	1	8
Testset2	22-Mar-18	100	8	5	3	6	2
Testset3	5-Apr-18	100	10	6	4	5	5
Testset4	19-Apr-18	100	8	5	3	4	4
Testset5	2-May-18	100	10	6	4	9	1
Total		**2751**	**114**	**61**	**53**	**83**	**31**

5.4 Functionality evaluation

To evaluate the ability of SeDan in discovery of plausible answers we were required to identify the questions that were initially unresolvable. To this end, during the first phase of the experiment, the original questions (i.e., SPARQL queries without any modification) were posed over the knowledge graph containing the semantics from Semantic Medline database, with the plausible reasoning *deactivated*. As Table 16 shows, only 52 questions[11] (45%), out of 114 questions, were answered using existing triples stored in the knowledge base. While 62 questions (55%), including 33 *causes* questions (out of 61 questions) and 29 *treatments* questions (out of 53 questions) were not resolvable (Table 17).

Table 16- Details of initially answered questions (without query modification)

Source	Total	Question domain		Question type	
		Causes	Treatments	Factoid	Yes/No
Train	34	20	14	33	1
Testset1	2	0	2	0	2
Testset2	2	1	1	2	0
Testset3	4	1	3	2	2
Testset4	4	2	2	3	1
Testset5	6	4	2	6	0
Total	**52**	**28**	**24**	**46**	**6**

[11] List of the initially answered queries can be found in Appendix II.

Table 17- Details of initially unanswered questions (without query modification)

Source	#Total	Question domain		Question type	
		Causes	Treatments	Factoid	Yes/No
Train	35	16	19	25	10
Testset1	7	3	4	1	6
Testset2	6	4	2	4	2
Testset3	6	5	1	3	3
Testset4	4	3	1	1	3
Testset5	4	2	2	3	1
Total	**62**	**33**	**29**	**37**	**25**

In the second phase of the experiment, the unanswered questions from the first phase were asked again, but this time with plausible reasoning *activated*. Table 18 shows SeDan was able to generate plausible queries for 42 out of 62 (68%) initially unanswered questions, and only 20 questions (32%) remained unresolvable[12] (Table 19).

A question is found unresolvable when the plausible reasoner is not able to generate any plausible queries for that question. Lack of relevant and supporting semantics in the knowledge graph, including SemMedDB, Drugbank and Disease Ontology, is the main reason of not being able to generate any plausible queries. This issue will be further discussed in the Discussion section.

Table 18- Details of questions with plausible queries

Source	Total	Question domain		Question type	
		Causes	Treatments	Factoid	Yes/No
Train	23	11	12	15	8
Testset1	7	3	4	1	6
Testset2	3	2	1	1	2
Testset3	5	5	0	2	3
Testset4	1	1	0	0	1
Testset5	3	1	2	1	2
Total	**42**	**23**	**19**	**20**	**22**

[12] List of plausibly answered questions and remained unresolvable questions can be found in appendices III and IV respectively.

Table 19- Details of remained unresolvable questions

Source	Total	Question domain		Question type	
		Causes	Treatments	Factoid	Yes/No
Train	12	5	7	10	2
Testset1	0	0	0	0	0
Testset2	3	2	1	3	0
Testset3	1	0	1	1	0
Testset4	3	2	1	1	2
Testset5	1	1	0	1	0
Total	**20**	**10**	**10**	**16**	**4**

Table 20 presents the plausibly answered questions and elaborates on the inputs of the plausible reasoning engine for each question, including:

- Type of question; depending on the answer that a question expects (i.e., *factoid* or *yes/no*), a SPARQL query can be a *Select* or *Ask*[13],
- Subsequent SPARQL queries using the semantics in the knowledge graph,
- Element to be substituted identifies the element (i.e., subject, object) of the SPARQL query that will be substituted in the process of plausible reasoning. For *Select* queries, in which one of the subject or object elements is a variable, the other non-variable element is the only option to be substituted. However, in *Ask* queries, subject, object or any combination of them could be replaced by their plausibly related semantics. Questions with decimal fractions show that the plausible reasoner investigated the substitution of both elements.
- Depth of plausibility identifies the length of the *plausible path* (**Definition 4**) in the process of plausible reasoning. In fact, depth of plausibility shows the number of hops from the *element to be substituted* that the query rewriting algorithm takes in order to discover the knowledge graph.

[13] Construct and Describe are two other types of SPARQL queries that are out of the scope of this experiment

For example, Question1 is looking for the *causes* of *Katayama fever*. The subsequent SPARQL query for this question will be:

```
SELECT ?x
WHERE { ?x   semp:CAUSES   obo:Katayama_fever. }
```

In the plausible resolution of this query, *obo:Katayama_fever*, the object element of the where statement, is the element of the query that will be substituted, and the query rewriting algorithm investigates up to a plausibility of depth of 4.

Likewise, Question 16 asks if *Saracatinib* was being considered as a treatment for *Alzheimer's disease*. The subsequent SPARQL query for this question will be:

```
ASK { semr:Antineoplastic_Agents   semp:TREATS   semr:Alzheimer's_Disease. }
```

In the case of Question 16, substituting one element (either subject or object) of the *where* statement is not effective in finding a plausible answer. In this regard, the plausible reasoner substitutes both subject (*semr:Antineoplastic*) and object (*semr:Alzheimer's_Disease*) with a *depth* of 2 for each element. In addition, insufficiency of substituting one element in the plausible resolution of Questions 20 and 37 compelled the plausible reasoner to substitute both the elements to find a plausible answer.

As seen in Table 20, some questions (e.g., Questions 9, 15, 27) have more than one entry, which are numbered with decimal fraction. In the case of *Ask* questions, in which both subject and object can be substituted, the plausible reasoner investigates the substitution of each element (if any plausible semantics exist).

Table 20- Details of the plausibly answered queries, the involved plausible patterns and the semantics conducting the patterns

#	BioASQ Question	SPARQL Query		Depth of Plausibility	Element to be substituted
		Type	**WHERE clause**		
1	What causes Katayama Fever?	SELECT	?x semp:CAUSES obo:Katayama_fever	4	Object
2	What is the cause of episodic ataxia type 6?	SELECT	?x semp:CAUSES obo:episodic_ataxia_type_6	5	Object
3	Do statins cause diabetes?	ASK	drugbank:Statins semp:CAUSES semr:Diabetes	4	Subject
4	Can Levoxyl (levothyroxine sodium) cause insomnia?	ASK	semr:Levothyroxine_Sodium semp:CAUSES semr:Primary_Insomnia	3	Object
5	Which antibodies cause Riedel thyroiditis?	SELECT	?x semp:CAUSES semr:Riedel's_thyroiditis	2	Object
6.1 6.2	Is the monoclonal antibody Trastuzumab (Herceptin) of potential use in the treatment of prostate cancer?	ASK	semr:Herceptin semp:TREATS semr:Prostate_cancer_metastatic	2 2	Object Subject
7	What is the treatment of acute myocarditis?	SELECT	?x semp:TREATS obo:acute_myocarditis	4	Object
8	What is the genetic basis of the Delayed Sleep-Phase Syndrome (DSPS)?	SELECT	?x semp:CAUSES semr:Delayed_Sleep_Phase_Syndrome	2	Object
9.1 9.2	Does DDX54 play a role in DNA damage response?	ASK	semr:DDX54 semp:CAUSES semr:DNA_Damage	2 3	Object Subject
10	Does a tonsillectomy affect the patient's voice?	ASK	semr:Secondary_post_tonsillectomy_hemorrhage semp:CAUSES semr:Voice_Disorders	3	Subject
11	Is there an RNAi drug being developed to treat amyloidosis?	ASK	semr:RNAIII semp:TREATS semr:Amyloidosis	3	Subject
12	Are there RNAi approaches considered for the treatment of kidney injury?	ASK	semr:RNAIII semp:TREATS semr:Injury_to_kidney_NOS	3	Subject
13	Has IVIG been tested in clinical trials for the treatment of Alzheimer's disease?	ASK	semr:Immunoglobulins,_Intravenous semp:TREATS semr:Alzheimer's_Disease	2	Object
14	Which bacteria causes erythrasma?	SELECT	?x semp:CAUSES semr:Erythrasma	2	Object
15.1 15.2	Do bacteria from the genus Morexella cause respiratory infections?	ASK	semr:Moraxella_Infections semp:CAUSES semr:Respiratory_Tract_Infections	2 2	Object Subject
16	Was Saracatinib being considered as a treatment for Alzheimer's disease?	ASK	semr:Antineoplastic_Agents semp:TREATS semr:Alzheimer's_Disease	{2, 2}	{Subject, Object}
17.1 17.2	Is celiac disease caused by gliadin-induced transglutaminase-2 (TG2)-dependent events ?	ASK	semr:transglutaminase_2 semp:CAUSES semr:Celiac_Disease	1 2	Object Subject
18.1 18.2	Can doxycycline cause photosensitivity?	ASK	semr:Doxycycline semp:CAUSES semr:Photosensitivity	2 2	Object Subject
19	What causes Black Lung?	SELECT	?x semp:CAUSES obo:black_lung	4	Object
20.1	Can canagliflozin cause euglycemic diabetic ketoacidosis?	ASK	drugbank:Canagliflozin semp:CAUSES semr:Diabetic_Ketoacidosis	{1, 5}	{Object, Subject}
20.2		ASK	semr:Glycosides semp:CAUSES semr:Euglycemic_Diabetic_Ketoacidosis	{3, 2}	{Subject, Object}

21	Mutation of which gene causes arterial tortuosity syndrome?	SELECT		2	Object
22	Can CD55 deficiency cause thrombosis?	ASK	semr:CD55 semp:CAUSES semr:Thrombosis	2	Subject
23	Which diseases are caused by mutations in Calsequestrin 2 (CASQ2) gene?	SELECT	semr:CASQ2 semp:CAUSES ?x	2	Subject
24	List disorders that are caused by mutations in the mitochondrial MTND6 gene.	SELECT	semr:Mitochondrial_DNA_mutation semp:CAUSES ?x	1	Subject
25	What organism causes woolsorter's disease	SELECT	?x semp:CAUSES obo:woolsorters_disease	4	Object
26	Which disease(s) are caused by HEX A deficiency?	SELECT	obo:hexosaminidase_A_deficiency semp:CAUSES ?x	4	Subject
27.1	Is Brucella abortus the organism that causes brucillosis known to cause spontaneous abortions in humans?	ASK	semr:Brucella_abortus_infection semp:CAUSES semr:Spontaneous_abortion	2	Object
27.2				2	Subject
28.1	Has rituximab been considered as a treatment for chronic fatigues syndrome?	ASK	semr:rituximab semp:TREATS semr:Chronic_Fatigue_Syndrome	2	Object
28.2				2	Subject
29	Dinutuximab is used for treatment of which disease?	SELECT	drugbank:Dinutuximab semp:TREATS ?x	3	Subject
30	What is the cause of Phthiriasis Palpebrarum?	SELECT	?x semp:CAUSES obo:Phthiriasis_Palpebrarum	4	Object
31	Orteronel was developed for treatment of which cancer?	SELECT	obo:Orteronel semp:TREATS ?x	4	Subject
32	Matuzumab has been tested for treatment of which cancers?	SELECT	drugbank:Matuzumab semp:TREATS ?x	3	Subject
33	Is nivolumab used for treatment of Non-Small-Cell Lung Cancer?	ASK	drugbank:Nivolumab semp:TREATS semr:Non-small_cell_lung_cancer_stage_II	3	Subject
34	Is lambrolizumab effective for treatment of patients with melanoma?	ASK	drugbank:Lambrolizumab semp:TREATS semr:melanoma	4	Subject
35	Which diseases can be treated with Afamelanotide?	SELECT	drugbank:Afamelanotide semp:TREATS ?x	3	Subject
36	List the diseases that can be treated using Vedolizumab.	SELECT	drugbank:Vedolizumab semp:TREATS ?x	3	Subject
37.1	Is Migalastat used for treatment of Fabry Disease?	ASK	drugbank:Migalastat semp:TREATS semr:Fabry_Disease	4	Subject
37.2			semr:Piperidines semp:TREATS semr:Fabry_Disease	{2, 2}	{Subject, Object}
38.2	Is ocrelizumab effective for treatment of multiple sclerosis?	ASK	drugbank:Ocrelizumab semp:TREATS semr:Multiple_Sclerosis	3	Subject
39	For the treatment of which conditions can atypical neuroleptic drugs be used?	SELECT	semr:Atypical_neuroleptic semp:TREATS ?x	3	Subject
40.1	Is tretinoin effective for photoaging?	ASK	semr:Tretinoin semp:TREATS semr:Photoaging	3	Object
41.1	Could Arimidex (anastrozole) cause hot flashes? (hot flushes)	ASK	semr:Arimidex semp:CAUSES semr:Hot_flushes	2	Subject
42	What is the definitive treatment for low pressure headache?	SELECT	?x semp:TREATS semr:Low_pressure_headache	3	Object

5.5 Correctness evaluation

When extending the knowledge coverage of medical knowledge-bases, the correctness of the plausibly inferred answers and the validity of the reasoning processes (i.e., the combination of the semantics supporting the query rewriting) are important as well. Hence, it is disproportionately important to find out (i) if the plausible reasoner can generate acceptable queries (Definition 9) from the clinical perspective (i.e., the plausible path and the supporting semantics reaching the plausible query are clinically acceptable), and (ii) the plausibly inferred answers are correct (Definition 8). The acceptability of the plausible queries was investigated by the domain expert. The correctness of the plausible answers was evaluated against the gold standard answers (released by BioASQ challenges) or verified by the domain expert.

Table 21 and Table 22 provide the details of SeDan's ability in generating acceptable plausible queries and finding the plausibly correct answers to the questions that were initially unresolvable. Table 21 lists the questions that expect factoid answers along with all the possible answers (a fact or a list of facts) for each question. For each question and its corresponding gold standard answer(s), the table identifies if (i) the knowledge graph (including SemMedDB, DrugBank and Disease ontology) includes the gold standard answer(s), (ii) SeDan generates an acceptable plausible query, and (iii) the set of plausible answers includes the gold standard answer(s).

Table 21- Ability of SeDan in finding the plausibly correct answer(s) for factoid questions – some questions have more than one gold standard answer that are listed as a number followed by a letter (x.a)

# Question	Gold standard answer	Acceptable plausible queries generated	Included in the KG	Matching answer was found
1	Schistosoma spp	✓	✓	✓
2	EAAT1 mutations	✓	O	O
5	IgG4	✓	✓	✓
7.a	Ibuprofen		✓	✓
7.b	Inotropic agents		✓	✓
7.c	Anti-inflammatory steroid and non-steroid drugs	✓	O	O
7.d	Mechanical support		O	O
8.a	Human leukocyte antigen allele		✓	✓
8.b	Human leukocyte antigen DRB1 allele		✓	✓
8.c	Human leukocyte antigen gene	✓	✓	✓
8.d	Circadian gene mutations		O	O
8.e	Structural polymorphisms in the hPer3		O	O
8.f	AA-NAT gene		O	O
14	Corynebacterium minutissimum	✓	✓	✓
19	Respirable coal mine dust	✓	✓	✓
21	SLC2A10 /GLUT10	✓	✓	✓
23.a	Catecholaminergic Polymorphic Ventricular Tachycardia	✓	✓	O
23.b	Familial Hypertrophic Cardiomyopathy		✓	O
23.c	Hypertrophic Cardiomyopathy		✓	✓
24.a	Leigh syndrome	✓	✓	✓
24.b	Leber's hereditary optic neuropathy		✓	✓
24.c	Dystonia		✓	✓
25	Bacillus Anthracis	✓	✓	✓
26.a	Tay-Sachs disease	✓	✓	✓
26.b	Chronic GM2 gangliosidoses		✓	✓
29	High-risk neuroblastoma	✓	✓	✓
30	Phthirus pubis	O	✓	O
31	Prostate cancer	✓	✓	✓
32.a	Pancreatic (Stage IV Pancreatic cancer)		✓	✓
32.b	Colorectal		✓	✓
32.c	Non-small cell lung		✓	✓
32.d	Ovarian (Ovarian cancer metastatic)		✓	✓
32.e	Pancreatic (Pancreatic carcinoma)	✓	✓	✓
32.f	Primary Peritoneal		✓	O
32.g	Gastric		✓	✓
32.h	Ovarian (Stage IC Ovarian cancer)		✓	O
32.i	Esophageal		✓	O
32.j	Cervical		✓	O
35.a	Erythropoietic Protoporphyria		✓	✓
35.b	Vitiligo		✓	✓
35.c	Hailey-Hailey disease	✓	✓	✓
35.d	Acne Vulgaris		✓	✓
35.e	Polymorphic light eruption		✓	✓
35.f	Actinic keratoses		✓	✓
36.a	Crohn's disease	✓	✓	✓
36.b	Ulcerative colitis		✓	✓
39.a	Schizophrenia		✓	✓
39.b	Schizoaffective disorder		✓	✓
39.c	Delusional disorder	✓	✓	✓
39.d	Psychotic relapse in neuroleptic malignant syndrome		✓	✓
39.e	Attention Deficit Hyperactivity Disorder		✓	✓
39.f	Psychotic disorders		✓	✓
42	Epidural blood patch	✓	O	O

For example, *Schistosoma spp* is the answer to Question 1, which is both included in the knowledge graph, and found by SeDan via a clinically acceptable plausible query. In the case of Question 2, SeDan generates acceptable plausible queries, but the knowledge graph does not include the gold standard answer (*EAAT1 mutations*) and consequently, SeDan is not able to find it—despite the fact that it can generate an acceptable plausible query.

Table 22- Ability of SeDan in finding the plausibly correct answer(s) for yes/no questions- some questions have been rewritten by transforming both subject and object of the statement that are listed as a number followed by a number (x.y)

# Question	**Gold standard answer**	**Acceptable plausible queries generated**	**Matching answer was found**	**# Question**	**Gold standard answer**	**Acceptable plausible queries generated**	**Matching answer was found**
3	Yes	✓	✓	20.1	Yes	✓	✓
4	Yes	O	-	20.2	Yes	✓	✓
6.1	Yes	✓	✓	22	Yes	✓	✓
6.2	Yes	✓	✓	27.1	Yes	✓	✓
9.2	Yes	✓	✓	27.2	Yes	✓	✓
10	Yes	✓	✓	28.1	Yes	✓	✓
11	Yes	✓	✓	28.2	Yes	✓	✓
12	Yes	✓	✓	33	Yes	O	O
13	Yes	✓	✓	34	Yes	✓	✓
15.1	Yes	O	-	37.1	Yes	✓	✓
15.2	Yes	✓	✓	37.2	Yes	✓	✓
16.1	Yes	✓	✓	38.1	Yes	✓	✓
17.1	Yes	O	-	38.2	Yes	✓	✓
17.2	Yes	✓	✓	40.1	Yes	✓	✓
18.1	Yes	✓	✓	41.1	Yes	✓	✓
18.2	Yes	✓	✓				

In Table 21, for those questions (i.e., questions 7, 8, 23, 39, etc.) that expect (list of) factoid answers, there are more than one possible answer listed. Each gold standard answer is numbered by the respective question's number, followed by a letter in the fraction.

Table 22 lists the questions that expect yes or no answers with their corresponding plausible resolutions. For each question and corresponding gold standard answer, the table identifies (i) if SeDan generates an acceptable plausible query, and (ii) if the set of plausible answers includes the correct answer(s). Among 22 questions expecting answer *Yes*, SeDan successfully finds the anticipated answer to all of them. But only 20 of the answers are acceptable, since for questions 4 and 33 the generated plausible queries were not approved by the expert.

Regarding the factoid questions, although in some cases, such as questions 7, 8 and 32, SeDan does not find all the possible gold standard answers, it can generate acceptable plausible queries for 19 (out of 20) questions, which 18 of them find the plausibly correct answers. Among the 53 possible gold standard answers for the 20 factoid questions, SeDan can find 44 of the answers. However, questions 2 and 42 remained unresolvable since the knowledge graph does not include their answers (i.e., *EAAT1 mutations* and *Epidural blood patch* respectively).

Overall, SeDan found the plausibly correct answers via acceptable plausible queries for 37 questions (out of 42) that were initially unanswered. Two of the questions (questions 2 and 42) are not answered since the knowledge graph does not include the correct answers. The other three questions (questions 4, 30, 33) are not resolved as their plausible queries were not clinically acceptable. Table 23 summarizes the results.

Table 23- Summary of SeDan's competence in finding acceptable plausible answers

Questions	Total	Question domain		Question type	
		Causes	Treatments	Factoid	Yes/No
Plausibly resolvable queries	42	23	19	20	22
Acceptable plausible queries (except questions 4, 30, 33)	39	21	18	19	20
With plausibly correct answers (except questions 2, 42)	37	19	18	18	19

In addition to the study of the correctness and acceptability of the plausible queries, it is worthwhile to investigate why some of the plausible queries are not acceptable. The domain expert was asked to identify the reason(s). Table 24 lists the clinically unacceptable plausible queries, the corresponding semantics in dispute, and the reason why a semantic is not acceptable.

For example, in the case of question 4, there are three semantics that were utilized through the plausible reasoning and their validity is in dispute:

Table 24- Unacceptable plausible queries and why the corresponding semantics are not acceptable

Question	Semantics in dispute	Issue
4	Ischemia ISA Sleep_disturbances	Wrong relationship
	Ischemia PRECEDES Shock	Vague/general concept
	Graves'_Disease ISA Mental_disorders	Wrong relationship
7.2	Infarction PRECEDES Myocarditis	Vague/general concept
8.3	Sleep_Disorders COEXISTS_WITH Celiac_Disease	Wrong relationship
	Syndrome COEXISTS_WITH Proteinuria	Vague/general concept
27.1	Spontaneous_abortion COEXISTS_WITH Arthritis	Wrong relationship
	Acute_infectious_disease PRECEDES Chronic_Disease	Vague/general concept
	Acute_infectious_disease COEXISTS_WITH Meningoencephalitis	Vague/general concept
30	Shock COEXISTS_WITH Infestation_by_Phthirus_pubis	Wrong relationship
	Hyperglycemia COEXISTS_WITH Infestation_by_Phthirus_pubis	Wrong relationship
	Infestation_by_Phthirus_pubis ISA Virus_Diseases	Wrong relationship
	Infestation_by_Phthirus_pubis COEXISTS_WITH Cerebrovascular_accident	Wrong relationship
32.6	inhibitors LOWER_THAN Peptides	Vague/general concept
	Peptides COEXISTS_WITH Paclitaxel	Wrong relationship
32.8	Bleomycin ISA Peptides	Wrong relationship
	Vinblastine COEXISTS_WITH Peptides	Wrong relationship
32.9	Cisplatin COEXISTS_WITH Peptides	Wrong relationship
32.10	Cisplatin COEXISTS_WITH Peptides	Wrong relationship
33	Peptides COEXISTS_WITH docetaxel	Wrong relationship

- First, from the clinical perspective, there is no relationship between *Ischemia* and *Sleep disturbances. Ischemia* is a restriction in blood supply tissue and is not a type of *Sleep disturbances*. Hence *Ischemia ISA Sleep_disturbances* is a *wrong relationship*.
- Second, *Ischemia* is a general term. The relationship *Ischemia PRECEDES Shock* does not identify that what kind of *Ischemia* (e.g., brain Ischemia, heart Ischemia, etc.) *precedes shock*. Although, some types of *Ischemia* (such as heart Ischemia) may lead to *shock*, there are some other types that do not cause a shock. Hence, the *ambiguity* of the semantic doesn't let the expert to verify the reasoning process.
- Third, the *Grave's Disease* is a thyroid disorder, and not a mental disorder. Thus, a *wrong relationship* is not acceptable for the domain expert.

Similarly, other clinically unacceptable plausible queries were declined due to either a wrong relationship (i.e., a relationship that does not make sense from the medical point of view), or a vague or a general concept that makes the relationship imprecise and not necessarily always true. These results raised the possibility that not all the semantics and relationships existing in SemMedDB are correct and clinically approved. This issue will be elaborated in the Discussion section.

5.6 Cost-effectiveness (performance) evaluation

To study the performance of SeDan and demonstrate its usability in real world practices, a set of measures are introduced to calculate the time required to perform different steps of the query rewriting algorithm:

- *SPARQL query execution time (SQET)* indicates the time required to ask a SPARQL query from the knowledge graph (RDF repository). Although this measure does not cover any steps of query rewriting algorithm and plausible reasoning, it provides a criterion to compare the cost of plausible reasoning with plain SPARQL query answering.
- *Retrieval time of applicable plausible semantics (ASRT)* measures the required time to perform the step 6 of Algorithm 1, in which the algorithm attempts to find a set of plausible semantics applicable to the body atom of the query (i.e., the element to be substituted) via the pattern matching function. The retrieval time is itemized by the plausible patterns.
- *Rewriting time of a plausible query (RWT)* measures the required time to perform the step 8 of Algorithm 1, in which one retrieved applicable semantics (from step 6) will be replaced with the to be substituted element in the query. This value shows

the time required to generate a plausible query from the initial query, or a previously generated plausible query.

- *Time to reach the first approved answer (FAAT)* identifies the time between the moment the user asks the question from SeDan and the moment SeDan finds the first plausibly correct answer (i.e., which is clinically acceptable as well)—i.e., it measures the total cost of finding the first acceptable plausible answer.

Table 25 summarizes the grand mean[14] of *SPARQL queries execution time (SQET)* for all the queries, broke down by the question domains and question types (the average SPARQL query execution time for each query is presented in Appendix V). As the table shows, it takes, on average, 5.6 milliseconds from SeDan to answer a SPARQL query. For the questions asking for treatments or the questions expecting fact(s) as their answers, the query execution time is higher than the average. Later, we will compare these values with the time required to perform different steps of the plausible reasoning and evaluate the usability of SeDan.

Table 25- Grand mean of execution time (ms) of SPARQL queries (summarized of values from Appendix V)

	Overall	**Question domain**		**Questions type**	
		Causes	**Treatments**	**Factoid**	**Yes/No**
Grand Mean (ms)	5.6	4.6	6.8	7.1	4.5

Table 26 recaps the average *retrieval time of applicable plausible semantics (ASRT)* (Appendix VI contains the detailed timings). ASRT computes the time that the pattern matching function searches the knowledge graph to find the applicable plausible semantics

[14] The mean of the means.

conforming with a specific plausible pattern. Table 26 provides an insight into the ASRT measure by breaking down the timings by sub-types of the plausible patterns.

The average retrieval time for all the plausible patterns falls in the (approximate) range of 5 to 19 milliseconds, except for the *a fortiori - more to less* pattern which is significantly higher. Our further investigation showed that there is an additional indexing/caching time in GraphDB that drastically increases the retrieval time of applicable semantics to the pattern *a fortiori-More to less* since it is the first pattern to investigate in the process of substituting a concept. Performing the experiment for few queries with a change in the order of the plausible patterns (e.g., by switching *generalization* to be the first pattern to be explored) proved the assumption.

Figure 11 depicts the average retrieval time of plausible semantics for each question by each plausible pattern. To keep the plot fathomable, the retrieval times of *a fortiori-more to less* and also retrieval times greater than 50 milliseconds (4 values) are eliminated.

Table 26- Grand mean of retrieval time of applicable plausible semantics (ms) to the concept to be substituted by plausible patterns (summary of values from Appendix VI)

	Overall time (ms)	Question domain		Questions type	
		Causes	Treatments	Factoid	Yes/No
A fortiori (More to less)	164.0	230.9	72.7	271.5	91.2
A fortiori (Less to more)	8.9	10.0	7.3	9.1	8.7
Generalization	12.0	13.2	10.3	11.9	12.0
Specification	9.5	9.2	10.0	8.0	10.5
Similarity (Hierarchical)	9.4	11.1	7.1	11.8	7.7
Similarity (Ordered-based)	17.2	18.5	15.4	16.1	17.9
Dissimilarity (Hierarchical)	7.6	8.7	6.1	7.1	7.9
Dissimilarity (Ordered)	6.4	6.4	6.4	5.1	7.2
Interpolation (Stands before)	14.4	16.5	11.5	14.7	14.2
Interpolation (Stands after)	9.4	6.4	13.6	5.5	12.1
Grand mean (A fortiori-More to less)	10.5	11.1	9.7	9.9	10.9

Although the values in Figure 11 are dispersed, conforming with Table 26, they show that the retrieval of the plausible semantics usually takes 5 to 20 milliseconds. In addition, the similar trend of the lines in the figure implies that retrieval time is consistent among different patterns for one question. In other words, if one question (e.g., Question 4) is costly with regard to a specific plausible pattern, it is possible that the retrieval of the semantics related to the other patterns is costly as well, and vice versa.

Regardless of the sudden rises and falls, over all, *similarity* and *interpolation-stands before* are the costliest patterns and *dissimilarity* and *interpolation-stands after* are the least. In addition to the complexity of the computation of each pattern in the pattern matching function, the number of applicable semantics to the concept in the question and, the number of relevant semantics to a plausible pattern in general, affect the required time to retrieve the applicable semantics.

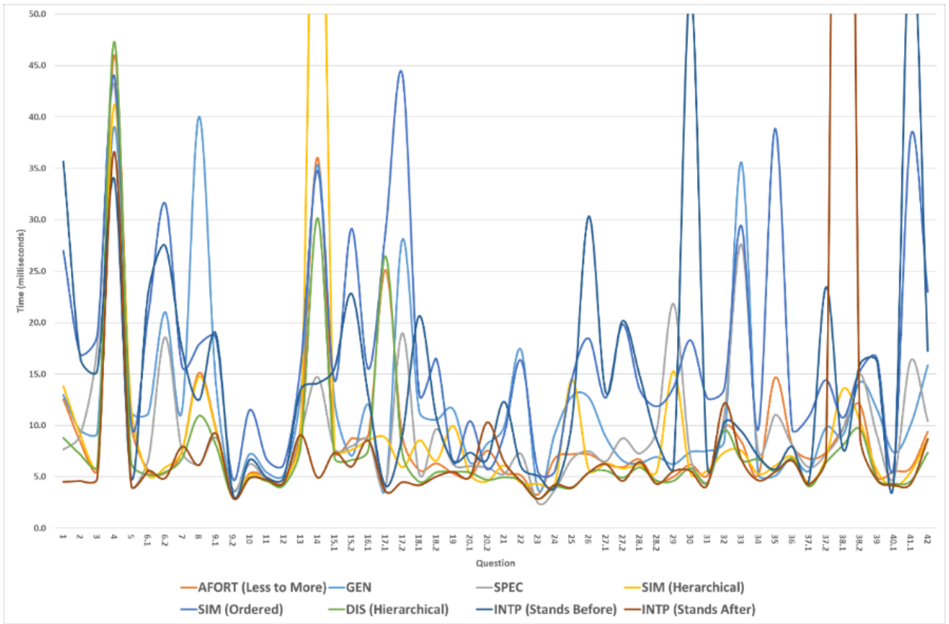

Figure 11- Average retrieval time (ms) of plausible semantics by plausible patterns for each question (values are from Appendix VI) – retrieval time of A fortiori-More to less and values bigger than 50 ms are cut off to keep the chart comprehensible

Table 27 provides the quantity of the existing semantics in the knowledge graph with regard to each plausible pattern. Although the numbers in the table represent the total number of the semantics (and not divided by each question or each concept), it shows that the number of plausible semantics applicable to the costly patterns is higher (Figure 11).

Table 27- Total number of relevant semantics to each pattern in the knowledge base

Plausible Patterns	Number of semantics
A fortiori	35648
Generalization	477158
Specialization	477158
Dissimilarity (Hierarchical)	5639
Similarity (Hierarchical)	41281
Similarity (Ordered)	1315959
Interpolation	682180

For example, Table 26 shows the highest retrieval time belongs to the similarity pattern, which has the biggest number of the applicable semantics in the knowledge graph (Table 27). Likewise, dissimilarity, which imposes the least overhead to the system, has the lowest number of the applicable semantics in the knowledge graph. Figure 12 depicts the trend between the number of applicable semantics to the plausible patterns and their correspoing retrieval time. The plot does not show a linear relationship between the number of the sematnics and the required time to retrieve them, but it does imply that a higher number of applicable semantics costs a greater retreival time from the system.

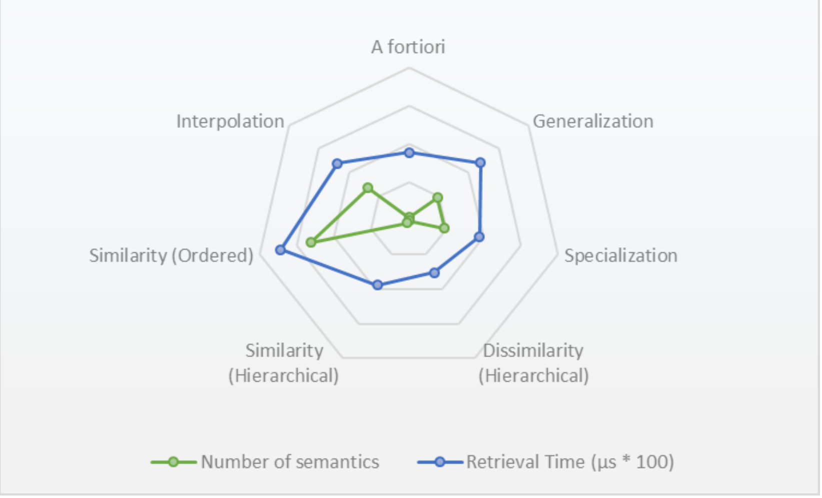

Figure 12- Comparison of the trends between the number of applicable semantics to plausible patterns and their retrieval time

Replacing the retrieved concept(s) and rewriting a query is the last overhead that the SeDan endures throughout the process of generating the plausible queries. Table 28 demonstrates the average *rewriting time of a plausible query (RWT)*; the time required to replace an element in the query (the step 8 of Algorithm 1)with a concept retrieved from previous step (the step 6 of Algorithm 1).

Table 28- Grand mean of the time (µs) required to generate a plausible query (summary of values from Appendix V)

	Overall	Question domain		Questions type	
		Causes	**Treatments**	**Factoid**	**Yes/No**
Grand Mean (µs)	20.4	15.2	27.4	16.0	23.5

As Table 28 shows, it approximately takes 20 microseconds to rewrite a SPARQL query by replacing an element (i.e., subject or object) with an applicable plausible semantic.

The measures above determine the time required to perform only *one* instance of the corresponding tasks. For example, *SPARQL query execution time* measures the time required to ask *one* SPARQL query*, retrieval time of applicable plausible semantics* identifies the cost of finding plausibly related semantics to the concept under investigation *and rewriting time of a plausible query* calculates the time of generating *one* plausible query. While, the process of performing plausible reasoning to find a plausible answer may contain hundreds (or thousands) of each of these steps.

Hence, in order to have a more realistic estimation of the cost of plausible reasoning, we introduced another measure, *time to reach the first approved answer*. This measure calculates the total time that SeDan takes to provide a plausibly correct answer to an initially unanswered question. This measure is comprised of the time required to execute the SPARQL queries (SQET), collect the applicable plausible semantics (ASRT) and rewrite the plausible queries (RWT):

$$\mathit{Time\ to\ reach\ the\ first\ approved\ answer} = \sum SQET + \sum ASRT + \sum RWT$$

Table 29 provides the information of the first approved plausible answers for each question. This information includes the time that SeDan needs to reach the answer, the plausibility depth that the answer is found at, and the plausible patterns that conducted the plausible reasoning.

For example, in the case of Question 1, it takes SeDan 8.2 seconds to reach the first approved answer in a plausibility depth of 3 (Table 29). Figure 13 depicts the reasoning process and the plausible paths that SeDan navigates to reach the node *Schistosoma spp* in the knowledge graph.

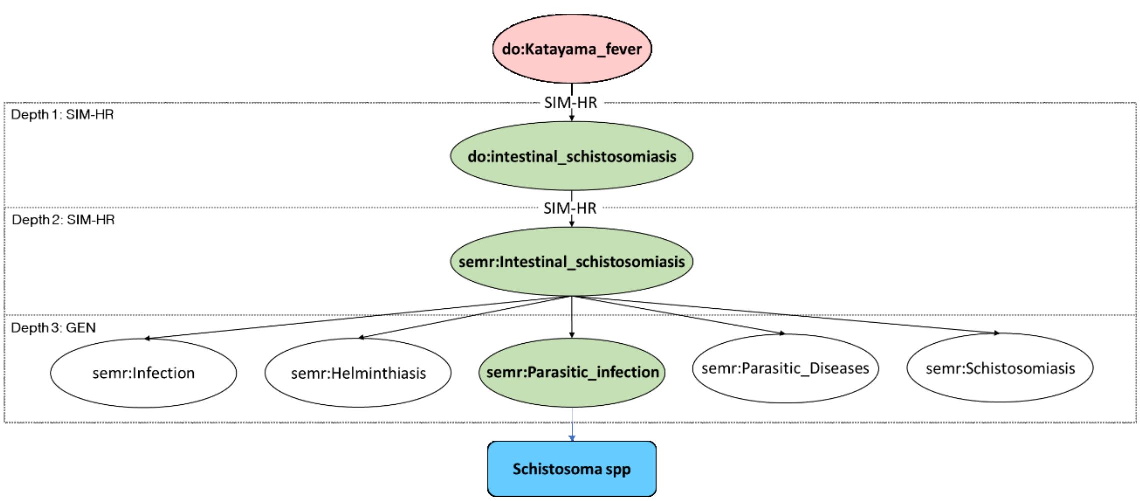

Figure 13- The process of plausible reasoning for Question 1

Table 29- Details of the first approved plausible answers for each question

#	Time (sec)	Depth	Plausible Patterns involved	#	Time (sec)	Depth	Plausible Patterns involved
1	8.2	3	{SIM-HR, SIM-HR, GEN}	26.a	0.6	3	{SIM-HR, SIM-HR, SIM-OR}
2	-	-	-	26.b	0.1	2	{SIM-HR, SIM-HR}
3	1.1	2	{SPEC, SIM-HR}	27.1	2.0	1	{SIM-OR}
4	-	-	-	27.2	1.1	1	{GEN}
5	1.6	1	{SIM-OR}	28.1	0.1	1	{AFORT-ML}
6.1	0.2	1	{GEN}	28.2	1.0	1	{AFORT-ML}
6.2	1.1	1	{GEN}	29	0.2	2	{GEN, SIM-HR}
7.a	17.1	4	{GEN, SIM-HR, AFORT-ML, SIM-OR}	30	-	-	-
7.b	-	-	-	31	2.4	4	{GEN, SIM-HR, GEN, GEN}
7.c	-	-	-	32.a	35.5	3	{GEN, SIM-HR, SIM-OR}
7.d	-	-	-	32.b	0.8	3	{GEN, SIM-HR, AFORT-ML}
8.a	3.5	2	{GEN, GEN}	32.c	4.6	3	{GEN, SIM-HR, GEN}
8.b	4.0	2	{GEN, SIM-OR}	32.d	3.3	3	{GEN, SIM-HR, GEN}
8.c	-	-	-	32.e	0.2	2	{GEN, SIM-HR}
8.d	-	-	-	32.f	-	-	-
8.e	-	-	-	32.g	4.6	3	{GEN, SIM-HR, GEN}
8.f	-	-	-	32.h	-	-	-
9.2	1.2	1	{SPEC}	32.i	-	-	-
10	5.4	2	{SIM-OR, SIM-OR}	32.j	-	-	-
11	9.9	2	{GEN, AFORT-LM}	33	-	-	-
12	7.5	2	{GEN, GEN}	34	0.2	3	{SIM-HR, GEN, SIM-HR}
13	0.3	1	{AFORT-LM}	35.a	18.2	3	{GEN, SIM-HR, SPEC}
14	24.4	2	{GEN, GEN}	35.b	5.6	3	{GEN, SIM-HR, GEN}
15.1	-	-	-	35.c	27.1	3	{GEN, SIM-HR, SIM-OR}
15.2	2.9	2	{SIM-OR, GEN}	35.d	3.1	3	{GEN, SIM-HR, AFORT-LM}
16.1	1.3	1	{AFORT-LM}	35.e	7.5	3	{GEN, SIM-HR, GEN}
17.1	-	-	-	35.f	7.3	3	{GEN, SIM-HR, GEN}
17.2	0.2	1	{GEN}	36.a	0.2	2	{GEN, SIM-HR}
18.1	0.2	1	{SPEC}	36.b	1.2	3	{GEN, SIM-HR, AFORT-LM}
18.2	1.3	1	{AFORT-ML}	37.1	0.2	3	{GEN, SIM-HR, SIM-OR}
19	0.3	2	{GEN, GEN}	37.2	5.1	2	{AFORT-ML, SIM-OR}
20.1	1.7	4	{GEN, GEN, SIM-HR, GEN}	38.1	1.3	2	{GEN, SIM-HR}
20.2	0.3	2	{GEN, GEN}	38.2	0.6	2	{GEN, SIM-HR}
21	13.0	2	{SIM-OR, SIM-OR}	39.a	0.1	1	{SIM-HR}
22	10.8	2	{SPEC, AFORT-LM}	39.b	0.1	1	{SIM-HR}
23.a	-	-	-	39.c	0.3	2	{SIM-HR, AFORT-ML}
23.b	-	-	-	39.d	0.1	1	{SIM-HR}
23.c	1.0	2	{SIM-OR, GEN}	39.e	0.3	2	{SIM-HR, AFORT-ML}
24.a	0.3	1	{SIM-OR}	39.f	0.3	2	{SIM-HR, AFORT-ML}
24.b	0.6	1	{SIM-OR}	40.1	0.1	2	{SIM-HR, SIM-OR}
24.c	0.9	1	{SIM-OR}	41.1	0.1	1	{SIM-HR}
25	123.2	4	{SIM-HR, SIM-HR, GEN, SIM-OR}	42	-	-	-

* GEN: Generalization, SPEC: Specialization, SIM: Similarity, HR: Hierarchical, OR: Ordered-based, AFORT: A Fortiori, ML: More to less, LM: Less to More, INTP: Interpolation, SB: Stands before, SA: Stands after

Figure 13 shows SeDan starts the process of plausible reasoning by replacing the *Katayama fever* with *intestinal schistosomiasis* (from the Disease ontology), which is retrieved via a *hierarchical similarity* relationship, *sameAs*. Since the plausibly generated query does not return any answers, SeDan moves further in the knowledge graph and replaces the concept with its equivalent in the SemMedDB, *semr:intestinal schistosomiasis*. Although this replacement leads to some plausibly inferred queries, none of them are acceptable. Then, SeDan proceeds to the plausible depth of 3, and finds five plausibly related semantics via a *generalization* relationship, *semp:isa*. By replacing the retrieved semantics, generating new plausible queries and asking the queries from the knowledge graph, the first plausibly correct answer will be found through *Parasitic_infection*.

In addition to the nodes that are illustrated in Figure 13, there are some middle steps that either returns no results or does not influence the reasoning process. However, these steps impose extra cost to the system. In total, this plausible navigation of the knowledge graph includes:

- 28 instances of SPARQL query execution
- 3 instances of plausible semantics retrieval, a fortiori – more to less
- 3 instances of plausible semantics retrieval, a fortiori – less to more
- 3 instances of plausible semantics retrieval, generalization
- 2 instances of plausible semantics retrieval, specialization
- 2 instances of plausible semantics retrieval, similarity – hierarchical
- 2 instances of plausible semantics retrieval, similarity – ordered-based
- 2 instances of plausible semantics retrieval, dissimilarity – hierarchical
- 2 instances of plausible semantics retrieval, dissimilarity – ordered-based
- 2 instances of plausible semantics retrieval, interpolation – stands before
- 2 instances of plausible semantics retrieval, interpolation – stands After

- 28 instances of plausible query generation

5.7 Effectiveness of the plausible patterns

To investigate the effectiveness and practicality of each plausible pattern in the process of plausible reasoning, the distribution of effective plausible patterns among the approved plausible answers is listed in Table 30. The table presents the number of clinically approved plausible queries to reach the plausibly correct answers, and the plausible patterns that are leveraged through the reasonings.

Table 30 shows that some patterns, like *generalization* and *similarity-order based* are used more frequently, and some other, like *interpolation* and *a fortiori*, are less frequent. Not all the plausible patterns are used equally in the reasoning processes, but all of them are helpful and necessary for a successful plausible reasoning.

Table 30- Distribution of effective plausible patterns among approved plausible answers

#	No. of clinically approved plausible reasoning	Effective Plausible Patterns						
		AFORT-ML	**AFORT-LM**	**GEN**	**SPEC**	**SIM-HR**	**SIM-OR**	**INTP**
1	2	-	-	✓	-	✓	-	-
3	6	✓	-	✓	✓	✓	-	✓
5	7	✓	✓	-	-	-	✓	-
6.1	2	-	✓	✓	-	-	-	-
6.2	2	-	-	✓	-	-	-	-
7.1	1	✓	-	✓	-	✓	✓	-
8.a	1	-	-	✓	-	-	-	-
8.b	1	-	-	✓	-	-	✓	-
9.2	4	-	-	✓	✓	-	-	✓
10	2	-	-	-	-	-	✓	✓
11	68	✓	✓	✓	✓	✓	✓	✓
12	43	✓	-	✓	✓	-	✓	✓
13	8	-	✓	✓		-	✓	-
14	5	-	-	✓	✓	-	✓	-
15.2	35	-	-	✓	-	-	✓	✓
16.1	3	-	✓	✓	-	-	✓	-
17.2	16	✓	✓	✓	-	-	-	✓

18.1	8	✓	-	✓	✓	-	✓	-
18.2	7	✓	✓	-	✓	-	✓	✓
19	9	-	-	✓	-	-	✓	-
20.1	5	-	✓	✓	-	✓	-	✓
20.2	2	-	-	✓	-	-	✓	-
21	1	-	-		-	-	✓	-
22	5	-	✓	✓	✓	-	✓	✓
23.c	1	-	-	✓	-	-	✓	-
24.a	7	-	-	-	-	-	✓	-
24.b	2	-	-	-	-	-	✓	-
24.c	5	-	-	-	-	-	✓	-
25	1	-	-	-	✓	✓	✓	-
26.a	3	-	-	-		✓	✓	-
26.b	2	-	-	✓	✓	✓	-	-
27.1	3	✓	✓	-	-	-	✓	✓
27.2	22	✓	-	✓	-	-	✓	✓
28.1	16	✓	✓	✓	-	-	-	✓
28.2	7	✓	✓	✓	-	-	✓	-
29	31	-	✓	✓	✓	✓	-	✓
31	2	-	✓	✓	-	✓	✓	-
32.a	1	-	-	✓	-	✓	✓	-
32.b	16	✓	-	✓	✓	✓	✓	✓
32.c	1	-	-	✓	-	✓	-	-
32.d	1	-	-	✓	-	✓	-	-
32.e	25	✓	✓	✓	-	✓	-	✓
32.g	2	-	-	✓	-	✓	✓	-
34	27	✓	✓	✓	✓	✓	-	✓
35.a	1	-	-	✓	✓	✓	-	-
35.b	19	-	-	✓	✓	✓	✓	-
35.c	1	-	-	✓	-	✓	✓	-
35.d	11	-	✓	✓	✓	✓	✓	-
35.e	1	-	-	✓	-	✓	-	-
35.f	2	-	-	✓	-	✓	✓	-
36.a	22	✓	✓	✓	-	✓	-	✓
36.b	14	✓	✓	✓	-	✓	-	✓
37.1	5	-	-	✓	✓	✓	✓	-
37.2	1	✓	-	-	-	-	✓	-
38.1	24	-	✓	✓	✓	✓	-	✓
38.2	1	-	-	✓	-	✓	-	-
39.a	26	✓	✓	✓	✓	✓	-	✓
39.b	26	✓	✓	✓	-	✓	-	✓
39.c	15	✓	✓	✓	-	✓	-	✓
39.d	29	✓	✓	✓	-	✓	✓	✓
39.e	8	✓	✓	✓	-	✓	-	✓
39.f	9	✓	✓	✓	✓	✓	-	-
40.1	10	✓	✓	✓	✓	✓	✓	✓
41.1	4	-	-	✓	✓	✓	-	✓

5.8 Discussion

Successful completion of the experiments illustrates that a plausible pattern, alone or in combination with other pattern(s), can provide plausible answer(s) and extends the query answering coverage of a knowledge graph. Aligned with three aspects of the evaluation framework, the designed experiment followed four main purposes to investigate the success of SeDan in achieving its objectives:

- to gain insight into the *functionality* and *behaviour* of the system,
- to report the level of success in providing a *complete* set of answers that are *plausibly correct*,
- to identify how *transparency* facilitates the understanding and acceptance of the reasoning processes,
- and ultimately, to recognize the areas to improve.

5.8.1 SeDan's potential use

SeDan is designed to help clinical care providers discover new knowledge and answer challenging questions that might not have been answered before. As mentioned earlier, for the time being, SeDan is considered as a query answering framework in research environments providing answers to medical researchers and scientists. It aims to provide insight from data for the domain experts who possess the domain knowledge to ask complex questions and validate the (plausible) results. We believe the semantic analysis of large data via plausible reasoning can provide meaningful insights by turning collected data into actionable knowledge. The acquired knowledge helps to better informed clinical diagnoses, improve personalized therapies, validate medical treatment and predict the adverse events to treatments, while lowering costs.

For example, precision medicine that proposes customized, target-based healthcare has recently become a popular research area. In line with the objectives of precision medicine, drug design tries to customize the drug delivery for each individual patient with the aim of maximizing therapeutic effects while minimizing undesired side effects (Wang et al., 2018). To this end, drug design studies the structure of (bio)molecules (such as protein) to understand the interactions between small organic molecule (i.e., ligand) and a receptor (i.e., patient) to suggest a therapeutic benefit for a specific patient (Klebe, 2013). Hence, precision medicine requires the integration of multiple data sources—such as biomedical, biological and biochemical data—to form a unified knowledge resource such as a knowledge graph to help experts find new knowledge and to answer specific questions. Given that an integrated knowledge resource may contain knowledge that is either not directly related to other knowledge elements (from another source) and the entire scope of the knowledge is not always known to the experts our plausible reasoning approach for semantic analytics, applied to the integrated knowledge graph and the associated data, is suitable to derive unknown and new knowledge relationships to advance the expert's understanding of biological structures and clinical processes in turn support the rather knowledge-intensive drug discovery progress.

Moreover, we believe that SeDan, offers a novel implementation of plausible reasoning over knowledge graphs—it is application is not limited to healthcare and medicine. We foresee the utility of SeDan for any problem domain (such as astrophysics, environment, and finance) that needs to derive new knowledge from large volumes of data and domain knowledge. In this regard, SeDan and the embedded plausible reasoning methods can provide new knowledge by associating knowledge atoms/elements to generate a plausible response or solution to a query—the plausible response is the first step in generating new knowledge (without recourse to evidence) which can subsequently be validated by domain experts to turn plausible to deterministic knowledge that can be used for generating actionable insights from semantic analytics. For example, SeDan can help data-driven recommender systems (e.g., movie recommendation systems) in two ways: (i) to provide deductive answer(s) to the questions that their answers are already included in the data, (ii) to leverage semantics, discover unknown, complex relationships, and suggest items that are explicable (i.e., SeDan can justify its findings) and may not be detectable by existing machine learning techniques.

5.8.2 Functionality

The functionality perspective of the evaluation is mainly focused on the proper development of SeDan, regarding the predefined characteristics of navigating the knowledge graph, exploring the applicable plausible semantics and ultimately generating the equivalent plausible queries. Hence, the functionality of SeDan could be respectively considered equal to the practicality of the query rewriting algorithm, the main module of the reasoning engine.

Table 18 showed that SeDan provides plausible resolutions for 42 out of 62 (68%) initially unanswered questions. Regardless of the acceptability of the plausible queries or correctness of the plausible answers, investigating the inference processes of the plausibly answered questions showed that SeDan navigates the knowledge graph and generates plausible queries as anticipated.

Starting with the concept in the question, it first explores the applicable semantics to the concept within the depth of plausibility of one (i.e., direct neighbours of the concept in the graph). The applicable semantics are explored based on the order of the plausible patterns (a fortiori, generalization, specialization, similarity, dissimilarity, interpolation), which is hard-coded in the algorithm—i.e., this sequence is not firm and could be modified.

Following the completion of visiting the concepts (i.e., nodes) within the depth of plausibility of one and retrieving the applicable semantics, the rewriting algorithm starts exploring the semantics within the depth of plausibility of two w.r.t. to the concept in the original question (i.e., or depth of plausibility of one w.r.t. to the concept in the generated plausible queries resulting from the first round of query rewriting). These iterations continue until the query rewriting algorithm visits all the applicable semantics within the depth of the plausibility identified by the user.

Code 8 reports the process of visiting an applicable plausible semantics at the plausibility depth of 3 and generating the subsequent plausible query. As seen, *Katayama fever*, the initial concept in the original question (Question 1), has been substituted by *Schistosomiasis* after three rounds of query rewiring: *Katayama fever* is first swapped by *intestinal schistosomiasis* from disease ontology via a *hierarchical similarity*. Then it moves to *SemMedDB* again via *hierarchical similarity* by replacing *obo:intestinal_schistosomiasis with semr:Intestinal_schistosomiasis*. Through the third rewriting, *Intestinal_schistosomiasis* will be substituted by its super-class, *semr:Schistosomiasis* (via generalization).

```
New Plausible Query: ?x semp:CAUSES semr:Schistosomiasis
Depth of Plausibility: 3
Sequence of Plausible Patterns: {SIM-HR, SIM-HR, GEN}
Supporting Semantics:
 {SIM-HR,(obo:intestinal_schistosomiasis oboInOwl:hasExactSynonym obo:Katayama_fever)}
 {SIM-HR,(semr:Intestinal_schistosomiasis owl:sameAs obo:intestinal_schistosomiasis)}
 {GEN,(semr:Intestinal_schistosomiasis semp:ISA semr:Schistosomiasis)}
```

Code 8. a generated plausible query during the exploration of the knowledge graph

Further investigation of the reasoning processes of the other plausibly inferred answers (like what is shown in Code 8) demonstrates that the query rewriting algorithm retrieves the correct plausible semantics, with regard to the concept to be substituted and the plausible pattern under investigation, and generates plausible queries precisely, as it is supposed to do.

It is worth emphasizing again that functionality evaluation does not consider the correctness of the plausible semantics leveraged during the reasoning nor the rationality of the sequence of the supporting semantics in the process of reasoning. It only confirms that SeDan, and to be specific, the query rewriting algorithm implements the plausible patterns flawlessly.

5.8.3 Comprehensiveness

Theoretically and practically (i.e., as functionality evaluation confirms) the plausible query rewriting algorithm is an exhaustive greedy search over the knowledge graph, which generates all the possible plausible queries within the desired depth of plausibility—i.e., if the time and resources allow.

However, the retrieval of the plausibly correct answer(s) strongly depends on the completeness of the semantics existing in the knowledge graph. The plausible query rewriting algorithm explores the plausible paths and generates plausible queries so long as an applicable plausible semantic is found—i.e., the query rewriting algorithm terminates at any point where the pattern matching function fails to retrieve an applicable semantics. The plausibly inferred answers are comprehensive to the same extent that the knowledge graph (as constituted by ontologies and instances) is complete. To put it differently, the completeness of the knowledge graph may have a spiral effect on the capability of the plausible reasoner to discover all the plausibly correct answers; a richer knowledge graph of semantics allows a more comprehensive set of plausible solutions.

In our experiments, we found out that providing additional semantics (e.g., a clinically acceptable semantic that is manually retrieved from the most recent released biomedical ontologies) could help the query rewriting algorithm to go further, generate acceptable plausible queries, and find plausible answers for 12 questions (out of 42 plausibly answered questions). In fact, we investigated how enrichment of the ontological semantics can improve the comprehensiveness of the answer set.

As a result, SeDan can plausibly answer 42 out of 62 (68%) initially unanswered questions, and only 20 questions (32%) remained unresolvable. These questions asked about a drug or disease that were neither included in the semantic knowledge graph nor helped by the manual insertion of additional semantics (if any was available). Hence, no exploration of the knowledge graph, albeit guided by the plausible patterns, can yield any plausible query. This observation verifies how the success of plausible reasoning depends on the richness and completeness of the available domain knowledge.

5.8.4 Plausible correctness

In any healthcare system it is essential to derive answers that don't lead to wrong diagnoses and treatments, and encourages the user confidence in the system (Clarke et al., 1994). As mentioned before, the correctness of the plausible answers has two facets: (i) if SeDan generates acceptable plausible queries (i.e., the plausible paths to the answers are clinically acceptable), and (ii) if the inferred answers are plausibly correct.

As Table 21 shows (summarized in Table 23) among 22 questions asking for yes/no answers, SeDan generates acceptable plausible queries for 20 questions, which for 19 of them the plausibly correct answers were found. In the case of the factoid questions, 18 questions (out of 20) were provided (at least one) plausibly correct answer—i.e., only one plausible query was not acceptable for the domain expert.

Overall, among all the 42 plausibly answered questions, the answers to 37 (88%) questions were plausibly correct and the corresponding plausible queries were clinically approved as well. These results show, despite the lack of support from crisp deductive reasoning, the plausible reasoning is able to provide valuable insights and make acceptable plausible inferences.

A reasonable and decent standard to measure the competence of SeDan in deriving the correct answers would be to compare its results with other participants of the BioASQ challenges that answered the same questions via different techniques. Therefore, we would be able to realize how our implementation of plausible reasoning over a health knowledge graph competes against other approaches (e.g., Natural Language Processing techniques, Machine Learning algorithms, etc.) in a practical setting. However, at this point, the details of the participants' performance in the previous BioASQ challenges are not available. Participating in the next BioASQ challenge is considered as future work.

5.8.5 Performance

As mentioned before, the response time of SeDan is comprised of: (i) retrieval of the applicable plausible semantics, (ii) generating plausible queries based on the retrieved semantics, and (iii) evaluating the resulting plausible SPARQL queries over the knowledge graph. Experimental results show, on average, SeDan spends 16.1 milliseconds to generate and evaluate a plausible query: 10.5 milliseconds to retrieve the applicable plausible semantics (Table 26), 20.4 microseconds to generate the corresponding plausible queries (Table 28) and 5.6 milliseconds to evaluate a query over the knowledge graph (Table 25).

To get a better insight of the performance of the plausible query rewriting algorithm, we compared its execution time with other well-known query rewriting algorithms proposed and designed for Description Logics (DLs) and OWL.

The CGLLR algorithm and the Requiem algorithm (Pérez-Urbina, Horrocks, & Motik, 2009) were evaluated over 9 different ontologies, using 5 questions per ontology. The ontologies vary significantly in size, regarding the numbers of classes (ranges between 2 to 194), properties (ranges between 1 to 31) and axioms (ranges between 2 to 222). Subsequently, the evaluation results show significant differences in the rewriting times of different questions in different ontologies—i.e., the number of queries generated in the rewriting ranges between 2 to 23,744 queries and the rewriting times range between 1 millisecond to 249 seconds, respectively.

Perez (Pérez-Urbina & Rodrıguez-Dıaz, 2012) evaluated the performance of Blackout, a highly optimized version of Requiem (Pérez-Urbina et al., 2010), by rewriting 14 LUBM queries with respect to four ABoxes with increasing size: 138K, 1.38M, 13.8M, and 138M triples. The results show growing rewriting times and evaluation times as the size of the ABoxes grows: from 26.3 milliseconds, in the smallest ABox, to 359 milliseconds in the biggest ABox—i.e., larger ABoxes exhaust up to 20 times more time on evaluating the rewritten queries.

Chortaras et. al. (Chortaras, Trivela, & Stamou, 2011) utilized the same data set as (Pérez-Urbina, Horrocks, et al., 2009) to evaluate their optimized query rewriting algorithm, Rapid. The results prove the efficiency of Rapid, in comparison to the original version of Requiem. However, the rewriting times range between 1 millisecond to 2.14 seconds.

The results above show the overall performance of a query rewriting algorithm strongly rely on the complexity (i.e., the number of classes, properties, etc.) of the ontologies that conduct the rewriting and the size of the (assertional) data that a rewritten query is submitted to. Comparing the results from SeDan with the other query rewriting algorithms shows the efficiency and applicability of the plausible query rewriting in real medical settings. However, like any other query rewriting algorithm, there still exist some potential to improve the performance of the plausible query rewriting.

The three elements of the response time depend on the complexity of the ontologies and the size of data. However, in the implementation of the plausible query rewriting, the order of the plausible patterns identifies the order of the applicable plausible semantics to be retrieved. This sequence determines the inference order of the plausible answers and, consequently, impacts the response time. In fact, the order of the plausible patterns in the pattern matching function does not directly impact the overall performance of the system, per se, but it profoundly influences the navigation behavior of the query rewriting algorithm over the knowledge graph.

The current sequence of the plausible patterns (hard-coded in the pattern matching function) is as *{AFORT-ML, AFORT-LM, GEN, SPEC, SIM-HR, SIM-OR, DISSIM-HR, DISSIM-OR, INTPA, INTPD}*. In the retrieval of applicable plausible semantics, the triples that match *a fortiori-more to less* pattern will be retrieved first, followed by the semantics that match *a fortiori-less to more* pattern, then the concepts that are connected via a *generalization* relationship, and so forth.

For example, it takes SeDan *17.1 seconds* to reach *Ibuprofen*, one of the plausible answers to Question 7, at the depth of 4 via the patterns *{GEN, SIM-HR, AFORT_ML, SIM-OR}*. With the current order of the patterns, at the plausibility depth of 1, the query rewriting algorithm first retrieves all the semantics matching with *a fortiori* before reaching the triple *obo:acute_myocarditis rdfs:subClassOf obo:myocarditis*, as a *generalization* association. Similarly, in the depth of 2, it investigates the associations conforming with *a fortiori*, *generalization*, and *specialization* patterns prior to reaching the *obo:myocarditis owl:sameAs semr:Myocarditis*, as a *similarity* association. Figure 14 depicts the engaged plausible patterns at each depth of plausibility to reach the plausible answer, Ibuprofen.

Depth 1	{**AFORT-ML, AFORT-LM, GEN,** *SPEC, SIM-HR, SIM-OR, DISSIM-HR, DISSIM-OR, INTPA, INTPD*}
Depth 2	{**AFORT-ML, AFORT-LM, GEN, SPEC, SIM-HR,** *SIM-OR, DISSIM-HR, DISSIM-OR, INTPA, INTPD*}
Depth 3	{**AFORT-ML,** *AFORT-LM, GEN, SPEC, SIM-HR, SIM-OR, DISSIM-HR, DISSIM-OR, INTPA, INTPD*}
Depth 4	{**AFORT-ML, AFORT-LM, GEN, SPEC, SIM-HR, SIM-OR,** *DISSIM-HR, DISSIM-OR, INTPA, INTPD*}

Figure 14- The involved plausible patterns at each depth of plausibility to reach Ibuprofen, as an answer to Question 7 – bold black font indicates the patterns involved during the reasoning process, and grey font shows the patterns that are not exploited. The underlined patterns show the exact sequence of the patterns that leads to the answer, Ibuprofen.

As Figure 14 implies any changes in the sequence of the patterns would impact the time to reach an answer. However, any increase or decrease in the response time resulting from a change in the sequence of the patterns is uncertain and hard to predict. That change strongly depends on the number of the applicable semantics to each plausible pattern regarding the concept under investigation (i.e., the first input of the pattern matching function, D, body atom of the query to change). Hence, finding a global optimal sequence of the plausible patterns is a complex task, which is considered as future work.

5.8.6 Scalability

The basic idea behind conventional query rewriting algorithms (Pérez-Urbina, Horrocks, et al., 2009; Rosati & Almatelli, 2010) is to exploit small schemas (ontological knowledge) to compile a modified query and evaluate it over large databases (instance data). Although, some optimized query rewriting algorithms cautiously incorporate ABox preprocessing to improve the performance of the algorithms (Kontchakov, Lutz, Toman, & Wolter, 2009), separating the domain knowledge from the instance data imposes some restrictions on perfect reformulation of the query—e.g., an ontology class with no assertions in the TBox (Rosati & Almatelli, 2010).

Today, the growing attention to ontologies by both the scientists and the industries has led to a development of expressive, huge ontologies with thousands of classes, properties and instances. In addition, the World Wide Web (WWW) hosts various autonomous and large data sources, which may be accompanied by their own ontologies or even, like SemMedDB, contain domain knowledge as well as millions of individual instances (Bennacer, Aufaure, Cullot, Sotnykova, & Vangenot, 2004). Hence, the separation of the knowledge/data sources may restrict the exploratory capabilities of the query rewriting.

An innovative setting that can efficiently handle the scalability of the computations would overcome the limitations of the conventional QR algorithms and to exploit the new opportunities that large-scale knowledge sources offer. Aligned with this objective, experimental results show SeDan has been a successful endeavor. The plausible query rewriting implements plausible patterns over a large knowledge graph (with over 11.5

million triple statements), which combines both ontological domain knowledge and assertional data. It derives plausible answers within the same response time span (in milliseconds to seconds) as its conventional predecessors, and yet guarantees the decidability. Although we have not performed the experiment with differently sized knowledge graphs and have not studied the performance of the system with smaller knowledge graphs, we anticipate that shrinking the knowledge graph will drop the number of the applicable plausible semantics, reduce the response time and adversely impact the correctness of the answers. However, it is an undeniable fact that the recent advances in the technology (e.g., processing power and memory capacity) have had a significant impact on this effort.

Furthermore, in ever-growing domains, such as healthcare and medicine, there are always new findings that are not included in the knowledge bases (or knowledge graphs). Hence, in any knowledge-based system, keeping the knowledge base updated with the most recent version of the sources (and new released sources) is a never-ending job. In addition to the scalability issues that continuous addition of new statements would cause, consistency maintenance of the knowledge graph will be a serious challenge.

5.8.7 Reliability

The in-depth experiment results show that response times from one question to another are widely scattered (from milliseconds to seconds). But it was also discussed that the SeDan's performance strongly depends on the number of the applicable semantics retrieved in each iteration of the query rewriting. Hence, the fluctuations in the response times are due to the different numbers of the retrieved semantics, which are expected and inevitable.

However, further investigation of the results shows among different runs of query rewritings of the same question, the response times vary slightly, while the plausible answers are derived consistently. Technological aspects (e.g., indexing or caching time of the repository, changes in the CPU usages, capacity of the memory) explains the changes in the response time. However, consistency in the order of the answers and plausible correctness of the results were anticipated as they directly depend on the order of the plausible patterns in the pattern matching function—i.e., consistent sequence of the plausible patterns guarantees the consistency in the order of the answers and plausible correctness of the answers.

5.8.8 Transparency

In the development and implementation of health care systems, human and organizational aspects are as important as technical issues (Yusof, Kuljis, Papazafeiropoulou, & Stergioulas, 2008). In healthcare analytics and medical query answering, the ultimate decisions will be taken by patients, physicians and health care providers. To trust a system, they require to understand how an outcome (e.g., diagnosis, treatment) is achieved. Hence a transparent, explicable analysis is paramount (Amarasingham, Patzer, Huesch, Nguyen, & Xie, 2014).

As discussed in the Results section (e.g., Code 8), SeDan facilitates and encourages transparency by providing the details of the query rewriting process: the generated plausible query, the sequence of the plausible patterns and the semantics driving the plausible exploration of the graph. This information permits the user to investigate the reasoning procedure in detail, evaluate the validity of each step, and finally accept or reject the derived plausible solution. Hence, we expect an effective clarification of the reasoning processes promotes health literacy, improves health communication, empowers user satisfaction, and ultimately user acceptance.

5.9 Summary

In this chapter we introduced the evaluation framework and its criteria to investigate the efficiency of SeDan and the plausible reasoner. The evaluation framework has three aspects: (i) *functionality* of the plausible reasoner, (ii) *correctness* of the plausible answers, and (iii) *cost-effectiveness* of the system. In the experiments, we leveraged the large-scale Semantic MEDLINE database, enriched with the standard clinical DrugBank and Disease ontologies, to answer intelligent medical questions from the BioASQ challenges.

Our real word experiment showed that even large knowledge sources (i.e., like SemMedDB with over 85 million records) may not be able to answer all the relevant questions as they usually suffer from incompleteness. The experimental results showed plausible reasoning, as an exploratory reasoning method, provides plausible resolutions for 42 out of 62 (68%) initially unanswered questions and expands the query answering coverage of the knowledge graph by 37 percent. It is important that 88 percent of the plausibly inferred answers and their corresponding reasoning processes (generated plausible queries) are clinically reasonable and acceptable for the domain expert.

Although the performance of the query rewriting algorithms varies among different domains and queries, the experimental results prove the efficiency and applicability of the plausible query rewriting in real medical settings.

Chapter 6: Conclusion and Future Work

In this thesis, we have implemented plausible reasoning over knowledge graphs, as a novel approach for performing semantics-based data analytics over large health data jointly with background knowledge, when available.

We recognized plausible reasoning as an exploratory reasoning method that leverages applicable semantic associations to perform inference through a set of frequently recurring patterns and suggest a plausible statement, which could be further tested deductively. To implement plausible reasoning, six well-known plausible patterns (generalization, specialization, similarity, dissimilarity, interpolation, and a fortiori) were identified.

Based on the semantic relationships exploited in the plausible patterns, we divided them into three main categories:

- Hierarchy-based patterns, including *generalization* and *specialization*, which navigate from a given set of objects to a larger (or smaller) set that contains (or is contained in) the given set;
- Ordered-based patterns, including *a fortiori* and *interpolation*, which leverage the partial order of the concepts with regard to a measurable feature to infer propositions that are implicit in the proposition with a higher degree of confidence;
- Equivalence (hybrid) patterns, including s*imilarity* and *dissimilarity*, which move between any concepts that are equal (or unequal) with regard to a hierarchical or order-based relationship.

We found that an effective representation of the plausible semantics (i.e., semantics that may conduct plausible patterns) is strongly dependent on fine-grained knowledge of how different concepts are semantically related. The Semantic Web framework provides formalisms to semantically represent data sources as knowledge graphs with various levels of expressivity (i.e., RDF(S), OWL), organized using domain-specific ontologies. Knowledge graphs encode semantic associations between different concepts and their properties, providing an opportunity to reason over the knowledge and reveal useful and unknown connections between the entities.

To implement the plausible patterns, we leveraged query rewriting as a Semantic Web querying technique that reformulates a given query to a modified version, which elicits both explicit (what a KB knows) and unknown (what it assumes) knowledge from the data. Within the Semantic Web framework, OWL 2 QL profile is designed to support a sound and complete query rewriting mechanism to answer queries through ontologies. OWL 2 QL is underpinned by the DL-Lite family of description logics. The Open World Assumption made in Description Logics makes OWL 2 QL suitable to work with incomplete knowledge in the Semantic Web scenarios.

Within the OWL 2 QL profile, the hierarchy-based patterns are supported via pre-existing hierarchical semantics (i.e., *rdfs:subClassOf* and *owl:instanceOf*), while the ordered-based patterns conduct plausible reasoning based on measurable relationships between concepts, such as size, chronological order, location, ranking or phase, which are not provided by the Semantic Web languages (such as OWL). To support the representation of the ordered-based plausible semantics and, consequently, facilitate the implementation of the plausible patterns, we introduced our plausible OWL extension (PL-OWL) to represent order-based semantics within the SW framework.

The introduced query rewriting algorithm performs as a graph traversal algorithm leveraging plausible patterns as its heuristics. The plausible query rewriting algorithm draws new associations, which were initially unknown, by conducting a pattern-driven exploration of semantic knowledge graphs to discover hidden associations. The SeDan framework combines all our efforts to implement plausible reasoning and manifests the semantics-based data analytics in one integrated system. SeDan is comprised of three main modules:

- *Plausible reasoner,* which the plausible query rewriting algorithm as its core component,
- *Knowledge sources*, which provide the data, semantics and ontological constructs needed to evaluate the queries and support the query rewriting process, and
- *User interface*, to accept the query, along with the desired plausible patterns, and communicate the plausible answer(s) and their justifications.

With the designed experiment we aimed to investigate (i) the functionality of the system, (ii) the validity of the results, and (iii) the performance of the system. The experimental evaluations showed that even large knowledge bases (e.g., SemMedDB with over 85 million records) suffer from incompleteness and may not be able to answer all the questions.

The results proved a plausible pattern, alone or in combination with other pattern(s), can discover complex associations and extend the query answering coverage of knowledge bases. In addition to the acceptable functionality and efficiency of our implementation of plausible reasoning over knowledge graphs, delivering the details of the reasoning process demonstrates our endeavour for a clear effective communication to improve user satisfaction and acceptance.

6.1 Limitation and Challenges

As mentioned earlier, the success of plausible reasoning depends on the richness and correctness of the data and domain knowledge that is captured and represented in the knowledge graph. Hence, keeping the knowledge graph updated with the latest available domain knowledge and including additional medical sources in the form of online medical Linked Data is a relentless challenge.

The experimental results raised some concerns regarding the correctness of the data. The validation of the plausible queries (by the expert) showed that (i) there are some associations that are clinically wrong, and (ii) one (or both) of the concepts in some associations are vague, which would not be generally acceptable as true associations. Obviously, the existing imprecision and inaccuracy in the data has influenced the outcome and acceptability of the plausibly inferred results. Hence, it is expected that a solid, clean, fully verified dataset will improve the soundness and comprehensiveness of the result set —i.e., a more complete and correct data/knowledge returns less unacceptable plausible queries and, probably, more clinically acceptable plausible answers.

To avoid extra complexity and to focus on the functionality and behaviour of the plausible patterns, the current implementation of the plausible query rewriting algorithm assumes (i) the initial query contains only one condition (i.e., one triple statement) in its WHERE clause, and (ii) only one component of the statement (subject or object) could be plausibly substituted in each round of plausible resolution of the query. This limitation didn't impact the designed experiment, since the retrieved questions from the BioASQ challenges are encoded to SPARQL queries with one condition in their WHERE clause. However, this limitation should be addressed in the improved version of the query rewriting algorithm.

Like any other health information and clinical decision support systems, the SeDan's cultural competence and acceptance within the environment that is going to be set up will be a challenge. As discussed earlier, we hope the transparency of the plausible reasoning could be of help.

6.2 Future Work

As briefly discussed in the results section, the present work can be extended in several directions:

- First, like any other query rewriting algorithm, the present version of the plausible query rewriting algorithm is amenable to optimization. The optimization will either improve the time taken to compute the rewritings or reduce the number of the plausible queries generated in the rewritings.

  One obvious optimization is to improve the order of the plausible patterns which retrieve the applicable plausible semantics. This optimization could help both aspects of the optimization. If the algorithm navigates the knowledge graph more purposefully, it won't explore ineffective semantics, will reduce the number of the generated plausible queries and will consequently reaches the plausible answer(s) faster. However, finding a global optimal sequence of the patterns is complicated. As the scattered distribution of effective plausible patterns among the approved plausible answers (Table 30) shows, finding the most practical patterns or ranking the patterns based on their applicability is not an easy task and requires more consideration.

Another optimization could be to prune the knowledge graph as the query rewriting navigates it. For example, in the case of *a fortiori*, two syntactically different semantics of *A standsBefore B* and *B standsAfter A* are semantically equivalent. Current implementation of the plausible query rewriting algorithm does not distinguish the equal semantics and entails them in the reasoning process. Although investigating the semantics at the running time may impose extra computation costs, it is worth studying.

- In the SPARQL query answering, users receive no answers because there is no statement (i.e., triple pattern) in the repository that matches with the conditions in the WHERE clause of the query. In SeDan, conforming with the Open World Assumption, we showed plausible reasoning is capable of deriving answers to the questions that were initially irresolvable. However, there are always some yes/no questions that the correct answer to them is *No*, for which providing a clinically reasonable explanation for that negative answer would be helpful for user satisfaction.

  Among the plausible patterns, *dissimilarity* has this potential to generalize the idea of *"if something is true about a concept or a phenomenon, it is plausibly not true for a concept or a phenomenon that is recognized as dissimilar"*. Hence, it would be a valuable addition if we could investigate the capability of SeDan to provide an explanation for *Ask* questions with *No* answers. Subsequently, we may be able to address *negation*, which was not at the focus of the study at this point.

- The experiments showed the injection of clinically correct semantics can initialize or resume plausible reasoning. It is also stated that keeping the knowledge graph updated with the most recent available knowledge sources is a continuous job.

However, providing this opportunity for the users (i.e., researchers or scientists) to manually insert statements and investigate their *hypotheses* (i.e., hypothesis testing) would be worthwhile.

As well as the inclusion of the expert's tacit knowledge in the form of individual facts, we previously showed (Mohammadhassanzadeh, Van Woensel, et al., 2017) analogical rules (i.e., plausible knowledge represented in the form of production rules) can derive new plausible facts/assertions and supplement deductive reasoning. Hence, including plausible rules would be a great addition to our plausible reasoning engine.

- Within the Semantic Web framework, rule languages such as SWRL and SPIN deliver deductive reasoning. In addition, OWL 2 RL, one of the OWL 2 profiles, is of interest in scalable reasoning without sacrificing too much expressivity. Simialr to *multi-strategy* reasoning systems (Woensel, Mohammadhassanzadeh, Abidi, & Abidi, 2015), SeDan could be augmented with other types of domain knowledge, such as deductive rules, and offer deductive reasoning along with plausible reasoning.
- In the experiments we leveraged all the available knowledge to evaluate the proficiency of SeDan in finding plausible answers to the questions that a large knowledge source, comprised of more than 11.5 million statements, could not provide any answers. Hence, we have no estimation how the plausible query rewriting algorithm performs with differently sized knowledge graphs. In this regard, evaluating the performance of SeDan with various sizes of knowledge graph is a part of future work.

- So far, six plausible patterns (generalization, specialization, similarity, dissimilarity, interpolation, and a fortiori) are formalized and included in the plausible query rewriting algorithm. However, based on our findings from argumentation studies on the crossroads of philosophy, reasoning, and logic, there are still more plausible patterns (i.e. apagoge, epagoge, etc.) that are worth studying (Aliseda, 2006; Hallaq, 2009).
- Participating in the next BioASQ challenges to find the competency of SeDan, and specifically the plausible reasoning engine in comparison to the other innovative techniques is one of the next steps.
- User satisfaction cannot be evaluated without the system being used by real users. In this regard, designing and conducting an experiment with the target users of the system, (e.g., physicians, medical researchers, scientists, etc.) is the next step.

Bibliography

Abell, A., Romero, O., Pedersen, T. B., Berlanga, R., Nebot, V., Aramburu, M. J., & Simitsis, A. (2017). Using semantic web technologies for exploratory OLAP: a survey. *IEEE Transactions on Knowledge and Data Engineering*, *27*(2), 571--588. Retrieved from http://ieeexplore.ieee.org/abstract/document/6834809/

Abidi, S. R., Cox, J., Abusharekh, A., Hashemian, N., & Abidi, S. S. R. (2016). A Digital Health System to Assist Family Physicians to Safely Prescribe NOAC Medications. *Studies in Health Technology and Informatics*, *228*, 519–23. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/27577437

Abidi, S. S. R. (2001). Knowledge management in healthcare: towards 'knowledge-driven' decision-support services. *International Journal of Medical Informatics*, *63*(1–2), 5–18. http://doi.org/10.1016/S1386-5056(01)00167-8

Abidi, S. S. R. (2007). Healthcare Knowledge Management: The Art of the Possible. In *Knowledge Management for Health Care Procedures* (pp. 1–20). Berlin, Heidelberg: Springer Berlin Heidelberg. http://doi.org/10.1007/978-3-540-78624-5_1

Abidi, S. S. R., & Abidi, S. (2013). An Ontology-Driven Personalization Framework for Designing Theory-Driven Self-management Interventions. In *Process Support and Knowledge Representation in Health Care* (pp. 97–112). Springer, Cham. http://doi.org/10.1007/978-3-319-03916-9_8

Abidi, S. S. R., & Hussain, S. (2007). Medical Knowledge Morphing via a Semantic Web Framework. In *Twentieth IEEE International Symposium on Computer-Based Medical Systems (CBMS'07)* (pp. 554–562). IEEE. http://doi.org/10.1109/CBMS.2007.70

Abidi, S. S. R., & Shayegani, S. (2009). Modeling the Form and Function of Clinical Practice Guidelines: An Ontological Model to Computerize Clinical Practice Guidelines. In *Workshop on Knowledge Management for Health Care Procedures* (pp. 81–91). Springer, Berlin, Heidelberg. http://doi.org/10.1007/978-3-642-03262-2_7

Abidi, S., Vallis, M., Piccinini-Vallis, H., Imran, S. A., & Abidi, S. S. R. (2018). Diabetes-Related Behavior Change Knowledge Transfer to Primary Care Practitioners and Patients: Implementation and Evaluation of a Digital Health Platform. *JMIR Medical Informatics*, *6*(2), e25. http://doi.org/10.2196/medinform.9629

Agibetov, A., Jiménez-Ruiz, E., & Solimando, A. (2015). Towards Shared Hypothesis Testing in the Biomedical Domain. In *Semantic Web applications and tools for life sciences (SWAT4LS)*. Cambridge.

Aleman-Meza, B., Halaschek-Wiener, C., Arpinar, I. B., Ramakrishnan, C., & Sheth, A. P. (2005). Ranking Complex Relationships on the Semantic Web. *IEEE Internet Computing*, *9*(3), 37–44. http://doi.org/10.1109/MIC.2005.63

Aleman-Meza, B., Sheth, A. P., Palaniswami, D., Eavenson, M., & Arpinar, I. B. (2006). Semantic Analytics in Intelligence: Applying Semantic. *Advanced Topics in Database Research*, *5*, 401.

Aliseda, A. (2006). Abductive Reasoning: Logical Investigations into Discovery and Explanation.

Almeida, R., Kaymak, U., & Sousa, J. (2010). A new approach to dealing with missing values in data-driven fuzzy modeling. *International Conference Fuzzy Systems*. Retrieved from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5584894

Amarasingham, R., Patzer, R. E., Huesch, M., Nguyen, N. Q., & Xie, B. (2014). Implementing Electronic Health Care Predictive Analytics: Considerations And Challenges. *Health Affairs*, *33*(7), 1148–1154. http://doi.org/10.1377/hlthaff.2014.0352

Anshakov, O., & Gergely, T. (2010). *Cognitive reasoning: A formal approach*. Springer Science & Business Media.

Antonelli, G. A. (2008). Non-monotonic logic. *Stanford Encyclopedia of Philosophy*.

Anyanwu, K., & Sheth, A. (2003). P-Queries: enabling querying for semantic associations on the semantic web. In *Proceedings of the twelfth international conference on World Wide Web - WWW '03* (p. 690). New York, New York, USA: ACM Press.

Ausín, D., López-de-Ipina, D., & Castanedo, F. (2014). A probabilistic OWL reasoner for intelligent environments. In *International Conference on Uncertainty Reasoning for the Semantic Web* (Vol. 1259).

Barcelona, A. (2012). *Metaphor and metonymy at the crossroads: A cognitive perspective*. Walter de Gruyter. Retrieved from https://books.google.ca/books?hl=en&lr=&id=aUogAAAAQBAJ&oi=fnd&pg=PP9&dq=metaphor+metonymy+plausible&ots=gsIMottHn4&sig=2CtYiONW8qlWD9rV4ikYIJHGlmg

Bennacer, N., Aufaure, M.-A., Cullot, N., Sotnykova, A., & Vangenot, C. (2004). Representing and Reasoning for Spatiotemporal Ontology Integration. In *OTM Confederated International Conferences" On the Move to Meaningful Internet Systems"* (pp. 30–31). http://doi.org/10.1007/978-3-540-30470-8_14

Berners-Lee, T., Hendler, J., & Lassila, O. (2001). The Semantic Web. A new form of Web content that is meaningful to computers will unleash a revolution of new possibilities. *Scientific American*, *284*(5), 1–5.

Berners-Lee, T., Hendler, J., & Lassila, O. (2001). The semantic web. *Scientific American*, *284*(5), 28–37.

Bianchi, F., Palmonari, M., Cremaschi, M., & Fersini, E. (2017). Actively Learning to Rank Semantic Associations for Personalized Contextual Exploration of Knowledge Graphs. In *European Semantic Web Conference* (pp. 120–135). Springer, Cham. http://doi.org/10.1007/978-3-319-58068-5_8

Bienvenu, M. (2016). Ontology-Mediated Query Answering: Harnessing Knowledge to Get More From Data. In *International Joint Conference on Artificial Intelligence*.

Billington, D. (2017). Principles and Examples of Plausible Reasoning and Propositional Plausible Logic. *ArXiv Preprint ArXiv:1703.01697*. Retrieved from https://arxiv.org/abs/1703.01697

Billington, D., Estivill-Castro, V., Hexel, R., & Rock, A. (2006). Using Temporal Consistency to Improve Robot Localisation. In *Robot Soccer World Cup* (pp. 232–244). Berlin, Heidelberg: Springer Berlin Heidelberg. http://doi.org/10.1007/978-3-540-74024-7_20

Bischof, S., Krötzsch, M., Polleres, A., & Rudolph, S. (2014). Schema-Agnostic Query Rewriting in SPARQL 1.1. In *International Semantic Web Conference* (pp. 584–600). http://doi.org/10.1007/978-3-319-11964-9_37

Blachowicz, J. (1989). Discovery and ampliative inference. *Philosophy of Science*, 438–462.

Bobillo, F., & Straccia, U. (2011). Fuzzy ontology representation using OWL 2. *International Journal of Approximate Reasoning*, *52*(7), 1073–1094.

Bouamrane, M. M., Rector, A., & Hurrell, M. (2011). Using OWL ontologies for adaptive patient information modelling and preoperative clinical decision support. *Knowledge and Information Systems*, *29*(2), 405–418. http://doi.org/10.1007/s10115-010-0351-7

Bouchon-Meunier, B., Dubois, D., Godo, L., & Prade, H. (1999). Fuzzy Sets and Possibility Theory in Approximate and Plausible Reasoning. In *Fuzzy sets in approximate reasoning and information systems* (pp. 15–190). Springer, Boston, MA. http://doi.org/10.1007/978-1-4615-5243-7_2

Bunnin, N., & Yu, J. (Eds.). (2004). *Plausibility*. *The Blackwell Dictionary of Western Philosophy*. Oxford, UK: Blackwell Publishing Ltd. http://doi.org/10.1111/b.9781405106795.2004.x

Calì, A., Frosini, R., Poulovassilis, A., & Wood, P. T. (2014). Flexible Querying for SPARQL. In *OTM Confederated International Conferences On the Move to Meaningful Internet Systems* (pp. 473–490). Springer, Berlin, Heidelberg. http://doi.org/10.1007/978-3-662-45563-0_28

Calvanese, D., Giacomo, G. De, & Lembo, D. (2007). Tractable reasoning and efficient query answering in description logics: The DL-Lite family. *Journal of Automated*. Retrieved from http://link.springer.com/article/10.1007/s10817-007-9078-x

Carvalho, R., Laskey, K., & Costa, P. (2013). PR-OWL 2.0–bridging the gap to OWL semantics. In *Uncertainty Reasoning for the Semantic Web II* (pp. 1–18). Springer.

Carvalho, R. N., Laskey, K. B., & Costa, P. C. G. (2017). PR-OWL – a language for defining probabilistic ontologies. *International Journal of Approximate Reasoning*, *91*, 56–79. http://doi.org/10.1016/J.IJAR.2017.08.011

Cellucci, C. (2013a). Philosophy of mathematics: Making a fresh start. *Studies in History and Philosophy of Science Part A*, *44*(1), 32–42. http://doi.org/10.1016/J.SHPSA.2012.09.002

Cellucci, C. (2013b). *Rethinking logic: Logic in relation to mathematics, evolution, and method*. *Springer*. Springer, Dordrecht. http://doi.org/10.1007/978-94-007-6091-2_18

Chen, W., Xiong, W., Yan, X., & Wang, W. (2018). Variational Knowledge Graph Reasoning. *ArXiv Preprint ArXiv:1803.06581*. Retrieved from http://arxiv.org/abs/1803.06581

Chiba, H., Nishide, H., & Uchiyama, I. (2015). Construction of an Ortholog Database Using the Semantic Web Technology for Integrative Analysis of Genomic Data. *PLOS ONE*, *10*(4), e0122802. http://doi.org/10.1371/journal.pone.0122802

Chortaras, A., Trivela, D., & Stamou, G. (2011). Optimized query rewriting for OWL 2 QL. *Automated Deduction–CADE-23*.

Clarke, K., O'Moore, R., Smeets, R., Talmon, J., Brender, J., McNair, P., … Barber, B. (1994). A methodology for evaluation of knowledge-based systems in medicine. *Artificial Intelligence in Medicine*, *6*(2), 107–121. http://doi.org/10.1016/0933-3657(94)90040-X

Cohen, G., & Conway, M. (2007). *Memory in the real world*.

Collins, A., & Michalski, R. (1989). The logic of plausible reasoning: A core theory. *Cognitive Science*, *13*, 1–49. http://doi.org/10.1016/0364-0213(89)90010-4

Consortium, W. W. W. (2014). RDF 1.1 Concepts and Abstract Syntax.

Costa, P. da, Laskey, K., & Laskey, K. (2008). PR-OWL: A Bayesian ontology language for the semantic web. In *Uncertainty Reasoning for the Semantic Web I* (pp. 88–107).

Costa, P. C. (2005). *Bayesian Semantics for the Semantic Web*. George Mason University.

Das, R., Dhuliawala, S., Zaheer, M., Vilnis, L., Durugkar, I., Krishnamurthy, A., … McCallum, A. (2017). Go for a Walk and Arrive at the Answer: Reasoning Over Paths in Knowledge Bases using Reinforcement Learning. *ArXiv Preprint ArXiv:1711.05851*. Retrieved from http://arxiv.org/abs/1711.05851

Davis, E. (1990). *Representations of commonsense knowledge*. Morgan Kaufmann.

Decker, S. (2007). *Detection of bursty and emerging trends towards identification of researchers at the early stage of trends*. South Carolina State University. Retrieved from https://getd.libs.uga.edu/pdfs/decker_sheron_l_200708_ms.pdf

Derrac, J., & Schockaert, S. (2014). Characterising semantic relatedness using interpretable directions in conceptual spaces. In *ECAI* (pp. 243--248).

Derrac, J., & Schockaert, S. (2015). Inducing semantic relations from conceptual spaces: A data-driven approach to plausible reasoning. *Artificial Intelligence*, *228*, 66–94. http://doi.org/10.1016/J.ARTINT.2015.07.002

Dezert, J. (2002). Foundations for a new theory of plausible and paradoxical reasoning. *Information and Security*, *9*, 13--57. Retrieved from http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.98.6027&rep=rep1&type=pdf

Dolog, P., Stuckenschmidt, H., Wache, H., & Diederich, J. (2009). Relaxing RDF queries based on user and domain preferences. *Journal of Intelligent Information Systems*, *33*(3), 239–260. http://doi.org/10.1007/s10844-008-0070-7

Dompere, K. (2012). *The Theory of the Knowledge Square: The Fuzzy Rational Foundations of the Knowledge-Production Systems*.

Dong, J. S., Feng, Y., Li, Y.-F., Tan, C. K.-Y., Wadhwa, B., & Wang, H. H. (2015). BOWL: augmenting the Semantic Web with beliefs. *Innovations in Systems and Software Engineering*, *11*(3), 203--215. Retrieved from https://idp.springer.com/authorize/casa?redirect_uri=https://link.springer.com/article/10.1007/s11334-015-0243-9&casa_token=WkHFs4eF6g4AAAAA:QXvCdL4Yq3cEfZVv3GTmhEboVm_ngjZiVDlJuFn_FgWB8xnWsOBq-nHtxaJF4EXSom6AdFK15cRbnZjqyA

Dontas, K., & Zemankova, M. (1990). APPLAUSE: An implementation of the Collins-Michalski theory of plausible reasoning. *Information Sciences*, *52*(2), 111–139. http://doi.org/10.1016/0020-0255(90)90038-C

dos Santos, L. L., Carvalho, R. N., Ladeira, M., Weigang, L., & Mendes, G. L. (2015). PR-OWL 2 RL-A Language for Scalable Uncertainty Reasoning on the Semantic Web. In *11th International Workshop on Uncertainty Reasoning for the Semantic Web (URSW 2015)* (p. 14).

Ehrlinger, L., & Wolfram, W. (2016). Towards a Definition of Knowledge Graphs. In *SEMANTiCS*. Retrieved from http://ceur-ws.org/Vol-1695/paper4.pdf

Elbassuoni, S., Ramanath, M., & Weikum, G. (2011). Query Relaxation for Entity-Relationship Search. In *Extended Semantic Web Conference* (pp. 62–76). Springer, Berlin, Heidelberg. http://doi.org/10.1007/978-3-642-21064-8_5

Endert, A. (2016). Semantic Interaction for Visual Analytics: Inferring Analytical Reasoning for Model Steering. *Synthesis Lectures on Visualization*, *4*(2), 1–99. http://doi.org/10.2200/S00730ED1V01Y201608VIS007

Fokou, G., Jean, S., & Hadjali, A. (2014). Endowing Semantic Query Languages with Advanced Relaxation Capabilities. In *International Symposium on Methodologies for Intelligent Systems* (pp. 512–517). Springer, Cham. http://doi.org/10.1007/978-3-319-08326-1_53

Fokou, G., Jean, S., Hadjali, A., & Baron, M. (2016). RDF Query Relaxation Strategies Based on Failure Causes. In *International Semantic Web Conference* (pp. 439–454). Springer, Cham. http://doi.org/10.1007/978-3-319-34129-3_27

Ge, S. X. (2017). Exploratory bioinformatics investigation reveals importance of junk DNA in early embryo development. *BMC Genomics*, *18*(1), 200. http://doi.org/10.1186/s12864-017-3566-0

Gil, Y. (2005). Description logics and planning. *AI Magazine*, *26*(2), 73. Retrieved from http://www.aaai.org/ojs/index.php/aimagazine/article/view/1814

Glimm, B., & Ogbuji, C. (2013). SPARQL 1.1 Entailment Regimes. W3C Recommendation, W3C, Mar. 2013. Retrieved October 31, 2018, from http://www.w3.org/TR/sparql11-entailment

Gnanambal, S. (2014). A framework for personalized decision support system for the healthcare application. *Health Science Journal*.

Grimm, S., & Motik, B. (2005). Closed World Reasoning in the Semantic Web through Epistemic Operators. *OWLED*.

Gunaratna, K. (2017). *Semantics-based Summarization of Entities in Knowledge Graphs*. Wright State University. Retrieved from http://rave.ohiolink.edu/etdc/view?acc_num=wright1496124815009777

Guo, Y., Pan, Z., & Heflin, J. (2005). LUBM: A benchmark for OWL knowledge base systems. *Web Semantics: Science, Services and Agents on the World Wide Web*, *3*(2), 158–182.

Habicht, J.-P., Victora, C., & Vaughan, J. P. (1999). Evaluation designs for adequacy, plausibility and probability of public health programme performance and impact. *International Journal of Epidemiology*, *28*(1), 10--18. Retrieved from https://academic.oup.com/ije/article-abstract/28/1/10/777525

Haider, N. Al, Abidi, S., Woensel, W. Van, & Abidi, S. S. (2014). Integrating existing large scale medical laboratory data into the semantic web framework. In *Big Data (Big Data), 2014 IEEE International Conference on*. Retrieved from http://ieeexplore.ieee.org/abstract/document/7004338/

Halford, G., Wilson, W., & Phillips, S. (2010). Relational knowledge: the foundation of higher cognition. *Trends in Cognitive Sciences*.

Hallaq, W. (2009). *Sharī'a: theory, practice, transformations*.

Han, P., Klein, W., & Arora, N. (2011). Varieties of uncertainty in health care a conceptual taxonomy. *Medical Decision Making*.

Harmelen, F. Van, & McGuinness, D. (2004). OWL web ontology language overview. *W3C Recommendation*.

Harris, S., & Seaborne, A. (2013). SPARQL 1.1 query language.

Heflin, J. (2001). *Towards the semantic Web: knowledge representation in a dynamic distributed environment*. *Submitted in partial fulfillment of the requirement for the degree of Doctor of Philosophy*. Universidade de Maryland em College Park.

Heit, E. (2000). Properties of inductive reasoning. *Psychonomic Bulletin & Review*, *7*(4), 569–592.

Hogan, A., Mellotte, M., Powell, G., & Stampouli, D. (2012). Towards Fuzzy Query-Relaxation for RDF. In *Extended Semantic Web Conference* (pp. 687–702). Springer, Berlin, Heidelberg. http://doi.org/10.1007/978-3-642-30284-8_53

Holder, L. B., Cook, D. J., & Djoko, S. (1994). Substucture Discovery in the SUBDUE System. In *KDD workshop* (pp. 169--180). Retrieved from http://www.aaai.org/Papers/Workshops/1994/WS-94-03/WS94-03-015.pdf

Hood, L., & Friend, S. H. (2011). Predictive, personalized, preventive, participatory (P4) cancer medicine. *Nature Reviews Clinical Oncology*, *8*(3), 184–187. http://doi.org/10.1038/nrclinonc.2010.227

Horrocks, I. (2003). From SHIQ and RDF to OWL: The making of a web ontology language. *Web Semantics: Science, Services and Agents on the World Wide Web*, *1*(1), 7–26.

Horrocks, I., & Patel-Schneider, P. (2004). SWRL: A semantic web rule language combining OWL and RuleML. *W3C Member Submission*, *21*, 79.

Horrocks, I., & Sattler, U. (2007). A Tableau Decision Procedure for SHOIQ. *Journal of Automated Reasoning*.

Horvitz, E., Heckerman, D., & Langlotz, C. (1986). A Framework for Comparing Alternative Formalisms for Plausible Reasoning. In *AAAI* (pp. 210--214). Retrieved from http://www.aaai.org/Papers/AAAI/1986/AAAI86-034.pdf

Hsiao, I.-H., Pandhalkudi Govindarajan, S. K., & Lin, Y.-L. (2016). Semantic visual analytics for today's programming courses. In *Proceedings of the Sixth International Conference on Learning Analytics & Knowledge - LAK '16* (pp. 48–53). New York, New York, USA: ACM Press. http://doi.org/10.1145/2883851.2883915

Huang, H., Liu, C., & Zhou, X. (2012). Approximating query answering on RDF databases. *World Wide Web*, *15*(1), 89–114. http://doi.org/10.1007/s11280-011-0131-7

Hurtado, C. A., Poulovassilis, A., & Wood, P. T. (2008). Query Relaxation in RDF. In *Journal on Data Semantics X* (pp. 31–61). Berlin, Heidelberg: Springer Berlin Heidelberg. http://doi.org/10.1007/978-3-540-77688-8_2

Inokuchi, A., Washio, T., & Motoda, H. (2000). An Apriori-Based Algorithm for Mining Frequent Substructures from Graph Data. In *European Conference on Principles of Data Mining and Knowledge Discovery* (pp. 13–23). http://doi.org/10.1007/3-540-45372-5_2

Ippoliti, E. (2008). Demonstrative and non-demonstrative reasoning by analogy. *Demonstrative and Non-Demonstrative Reasoning in Mathematics and Natural Science*, 309–338.

Jafarpour, B., Abidi, S. R., & Abidi, S. S. R. (2016). Exploiting Semantic Web Technologies to Develop OWL-Based Clinical Practice Guideline Execution Engines. *IEEE Journal of Biomedical and Health Informatics*, *20*(1), 388–398. http://doi.org/10.1109/JBHI.2014.2383840

Jafarpour, B., Raza Abidi, S., Van Woensel, W., & Raza Abidi, S. S. (2019). Execution-time integration of clinical practice guidelines to provide decision support for comorbid conditions. *Artificial Intelligence in Medicine*, *94*, 117–137. http://doi.org/10.1016/J.ARTMED.2019.02.003

Jain, S., Gupta, C., & Bhardwaj, A. (2018). Research Directions Under the Parasol of Ontology Based Semantic Web Structure. In *International Conference on Soft Computing and Pattern Recognition* (pp. 644–655). http://doi.org/10.1007/978-3-319-60618-7_63

Jaynes, E. T. (2003). *Probability Theory: The Logic of Science*. Cambridge university press. Retrieved from https://books.google.ca/books?hl=en&lr=&id=UjsgAwAAQBAJ&oi=fnd&pg=PR17&dq=Probability+Theory:+The+Logic+of+Science&ots=OUNsgS-oex&sig=adgQh0mWQwi_S-l5YvHORV1VyBI#v=onepage&q=Probability Theory%3A The Logic of Science&f=false

Jing, Y., Jeong, D., & Baik, D.-K. (2009). SPARQL graph pattern rewriting for OWL-DL inference queries. *Knowledge and Information Systems*, *20*(2), 243--262. Retrieved from https://idp.springer.com/authorize/casa?redirect_uri=https://link.springer.com/article/10.1007/s10115-008-0169-8&casa_token=JsrxWy2MFa4AAAAA:YOjackx9mOKYIkUD8crqknZURFiIFd0L7crEXmYH93B7r6dGrTQINNGDVmV-Iz7xC095XM3tu-vRz4HImA

Jøsang, A. (1997). Artificial Reasoning with Subjective Logic. *Proceedings of the Second Australian Workshop on Commonsense Reasoning*, *48*, 34.

Kaisler, S. H., Espinosa, J. A., Armour, F., & Money, W. H. (2014). Advanced Analytics--Issues and Challenges in a Global Environment. In *ieeexplore.ieee.org* (pp. 729--738). Retrieved from https://ieeexplore.ieee.org/abstract/document/6758694/

Kell, D. B., & Oliver, S. G. (2004). Here is the evidence, now what is the hypothesis? The complementary roles of inductive and hypothesis-driven science in the post-genomic era. *Bioessays*, *26*(1), 99--105. Retrieved from http://onlinelibrary.wiley.com/doi/10.1002/bies.10385/full

Kharlamov, E., Jimnez-Ruiz, E., Zheleznyakov, D., Bilidas, D., Giese, M., Haase, P., … Koubarakis, M. (2013). Optique: Towards OBDA systems for industry. In *Extended Semantic Web Conference* (pp. 125--140).

Khoo, C. S., & Na, J.-C. (2006). Semantic relations in information science. *Annual Review of Information Science and Technology*, *40*(1), 157--228. Retrieved from https://dr.ntu.edu.sg/handle/10220/18367

Kibbe, W., Arze, C., Felix, V., & Mitraka, E. (2015). Disease Ontology 2015 update: an expanded and updated database of human diseases for linking biomedical knowledge through disease data. *Nucleic Acids Research*, *43*(D1), D1071–D1078. Retrieved from https://academic.oup.com/nar/article-abstract/43/D1/D1071/2435381

Kilicoglu, H., Shin, D., Fiszman, M., Rosemblat, G., & Rindflesch, T. C. (2012). SemMedDB: a PubMed-scale repository of biomedical semantic predications. *Bioinformatics*, *28*(23), 3158–3160. http://doi.org/10.1093/bioinformatics/bts591

Kim, J., Ostrowski, D. A., Yamaguchi, H., & Sheu, P. C.-Y. (2013). Semantic computing and business intelligence. *International Journal of Semantic Computing*, *07*(01), 87–117. http://doi.org/10.1142/S1793351X13500013

Klebe, G. (2013). Three-Dimensional Structure of Biomolecules. In *Drug Design* (pp. 291–314). Berlin, Heidelberg: Springer Berlin Heidelberg. http://doi.org/10.1007/978-3-642-17907-5_14

Klinov, P., & Parsia, B. (2008). Pronto: Probabilistic Ontological Modeling in the Semantic Web. In *Proceedings of the 2007 International Conference on Posters and Demonstrations-Volume 401* (pp. 82–83).

Knox, C., Law, V., Jewison, T., Liu, P., Ly, S., Frolkis, A., … Wishart, D. S. (2011). DrugBank 3.0: a comprehensive resource for “Omics” research on drugs. *Nucleic Acids Research*, *39*(Database), D1035–D1041. http://doi.org/10.1093/nar/gkq1126

Kolovski, V., Wu, Z., & Eadon, G. (2010). Optimizing enterprise-scale OWL 2 RL reasoning in a relational database system. In *International Semantic Web Conference*. Retrieved from https://link.springer.com/chapter/10.1007/978-3-642-17746-0_28

Kontchakov, R., Lutz, C., Toman, D., & Wolter, F. (2009). Combined FO rewritability for conjunctive query answering in DL-Lite. In *CEUR Workshop Proceedings*. Retrieved from http://eprints.bbk.ac.uk/1225/1/1225.pdf

Kostylev, E. V, & Grau, B. C. (2015). Semantics of SPARQL under OWL 2 Entailment Regimes. In *Description Logics*.

Krotzsch, M. (2017). Ontologies for knowledge graphs. In *Proc. 30th Int. Workshop on Description Logics (DL'17). CEUR Workshop Proceedings* (Vol. 1879). Retrieved from https://iccl.inf.tu-dresden.de/w/images/4/4c/Ontologies-for-knowledge-graphs-kroetzsch-DL2017.pdf

Krumholz, H. M. (2014). Big Data And New Knowledge In Medicine: The Thinking, Training, And Tools Needed For A Learning Health System. *Health Affairs*, *33*(7), 1163–1170. http://doi.org/10.1377/hlthaff.2014.0053

Kuipers, B. (1979). On representing commonsense knowledge. *Associative Networks: The Representation and Use of Knowledge by Computers*, 393–408.

Laborie, S., Ravat, F., Song, J., & Teste, O. (2015). Combining business intelligence with semantic web: Overview and challenges. *Actes Du XXXIIIème Congrès INFORSID, Biarritz*, 99–114.

Laskey, K. J., Laskey, K. B., Costa, P. C., Kokar, M. M., Martin, T., & Lukasiewicz, T. (2008). Uncertainty reasoning for the world wide web. *W3c Incubator Group Report*, *31*.

Law, V., Knox, C., Djoumbou, Y., & Jewison, T. (2014). DrugBank 4.0: shedding new light on drug metabolism. *Nucleic Acids Research*, *42*(D1), D1091–D1097. Retrieved from https://academic.oup.com/nar/article-abstract/42/D1/D1091/1047263

Lin, H. K., Harding, J. A., & Tsai, W. C. (2012). A rule-based knowledge system on semantic web for collaboration moderator services. *International Journal of Production Research*, *50*(3), 805–816. http://doi.org/10.1080/00207543.2010.544074

Liu, F., Huang, Y., & Lin, Y. (2013). Fuzzy ontology respresentation with fuzzy owl. In *2013 International Conference on Machine Learning and Cybernetics* (Vol. 4, pp. 1752–1757).

Lukasiewicz, T. (2017). Uncertainty Reasoning for the Semantic Web. In *Reasoning Web International Summer School* (pp. 276–291). Springer, Cham. http://doi.org/10.1007/978-3-319-61033-7_8

Mangold, C. (2007). A survey and classification of semantic search approaches. *International Journal of Metadata, Semantics and Ontologies*, *2*(1), 23. http://doi.org/10.1504/IJMSO.2007.015073

Martin, M., Weibel, R., Rocke, C., & Boker, S. M. (2018). Semantic Activity Analytics for Healthy Aging: Challenges and Opportunities. *IEEE Pervasive Computing*, *17*(3), 73–77. http://doi.org/10.1109/MPRV.2018.03367738

Matheus, C. J., Baclawski, K., Kokar, M. M., Letkowski, J. J., Matheus, C. J., Baclawski, K., … Letkowski, J. J. (2005). Using SWRL and OWL to capture domain knowledge for a situation awareness application applied to a supply logistics scenario. In *International Workshop on Rules and Rule Markup Languages for the Semantic Web* (pp. 130--144). Retrieved from https://link.springer.com/chapter/10.1007/11580072_11

Mathew, P. S., & Pillai, A. S. (2016). Big Data Challenges and Solutions in Healthcare: A Survey. In *Innovations in Bio-Inspired Computing and Applications* (pp. 543–553). Springer. http://doi.org/10.1007/978-3-319-28031-8_48

Mehdi, G., Brandt, S., Roshchin, M., & Runkler, T. A. (2016). Semantic Framework for Industrial Analytics and Diagnostics. In *IJCAI* (pp. 4016--4017).

Miller, J. A., Ramaswamy, L., Kochut, K. J., & Fard, A. (2015). Research directions for big data graph analytics. In *Big Data (BigData Congress), 2015 IEEE International Congress on* (pp. 785--794). Retrieved from http://ieeexplore.ieee.org/abstract/document/7207314/

Milstein, B., & Wetterhall, S. (1999). Framework for program evaluation in public health. *MMWR Recommendations & Reports*. Retrieved from https://stacks.cdc.gov/view/cdc/5204

Mohammadhassanzadeh, H., Abidi, S. R., Van Woensel, W., & Abidi, S. S. R. (2018). Investigating Plausible Reasoning Over Knowledge Graphs for Semantics-Based Health Data Analytics. In *IEEE 27th International Conference on Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE)* (pp. 148–153).

Mohammadhassanzadeh, H., Raza Abidi, S., Shah, M. S., Karamollahi, M., & Abidi, S. S. R. (2017). SeDAn: a Plausible Reasoning Approach for Semantics-based Data Analytics in Healthcare. In *WAIAH 2017, held within the AI*IA 2017 Conference*. Bari, Italy.

Mohammadhassanzadeh, H., Van Woensel, W., Abidi, S. R., & Abidi, S. S. R. (2017). Semantics-based plausible reasoning to extend the knowledge coverage of medical knowledge bases for improved clinical decision support. *BioData Mining*, *10*(1), 7. http://doi.org/10.1186/s13040-017-0123-y

Mohammadhassanzadeh, H., Woensel, W. Van, Abidi, S., & Abidi, S. S. R. (2016). A Semantic Web-based Approach to Plausible Reasoning for Improving Clinical Knowledge Engineering. In *BHI`16*.

Morgenthaler, S. (2009). Exploratory data analysis. *Wiley Interdisciplinary Reviews: Computational Statistics*, *1*(1), 33–44. http://doi.org/10.1002/wics.2

Nijssen, S., & Kok, J. N. (2004). A quickstart in frequent structure mining can make a difference. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining* (pp. 647--652). Retrieved from https://dl.acm.org/citation.cfm?id=1014134

Nikolaou, C., & Koubarakis, M. (2016). Querying incomplete information in RDF with SPARQL. *Artificial Intelligence*, *237*, 138--171. Retrieved from https://www.sciencedirect.com/science/article/pii/S0004370216300467

NISO Standard (ANSI). Guidelines for the Construction, Format, and Management of Monolingual Controlled Vocabularies (2010).

Nute, D. (2001). Defeasible logic. In *International Conference on Applications of Prolog* (pp. 151–169).

Ogiela, L. (2013). Semantic analysis and biological modelling in selected classes of cognitive information systems. *Mathematical and Computer Modelling*, *58*(5–6), 1405–1414. http://doi.org/10.1016/J.MCM.2012.12.001

Oroumchian, F., & Oddy, R. N. (1996). An application of plausible reasoning to information retrieval. In *19th international ACM conference on Research and development in information retrieval - SIGIR '96* (pp. 244–252). New York, New York, USA: ACM Press. http://doi.org/10.1145/243199.243271

Panahiazar, M., Taslimitehrani, V., Jadhav, A., & Pathak, J. (2014). Empowering personalized medicine with big data and semantic web technology: promises, challenges, and use cases. In *Big Data (Big Data), 2014 IEEE International Conference on* (pp. 790--795). Retrieved from http://ieeexplore.ieee.org/abstract/document/7004307/

Panton, K., Matuszek, C., & Lenat, D. (2006). Common sense reasoning–from Cyc to intelligent assistant. *Ambient Intelligence in …*.

Paulheim, H. (2017). Knowledge graph refinement: A survey of approaches and evaluation methods. *Semantic Web*, *8*(3), 489--508. Retrieved from https://content.iospress.com/articles/semantic-web/sw218

Peirce, C. S. (Charles S., & Ketner, K. L. (1992). *Reasoning and the logic of things : the Cambridge conferences lectures of 1898*. Harvard University Press. Retrieved from https://books.google.ca/books?hl=en&lr=&id=BWfRF5BAcG0C&oi=fnd&pg=PR11&dq=Reasoning+and+the+logic+of+things&ots=Zskw5GVs40&sig=5fZDTbWve7cK6Frgdi9H4ROleW8#v=onepage&q=Reasoning and the logic of things&f=false

Pérez-Urbina, H., Horrocks, I., & Motik, B. (2009). *Efficient query answering for OWL 2*.

Pérez-Urbina, H., Motik, B., & Horrocks, I. (2009). A Comparison of Query Rewriting Techniques for DL-lite. *Description Logics*.

Pérez-Urbina, H., Motik, B., & Horrocks, I. (2010). Tractable query answering and rewriting under description logic constraints. *Journal of Applied Logic*.

Pérez-Urbina, H., & Rodrıguez-Dıaz, E. (2012). Evaluation of query rewriting approaches for OWL 2. *Proc. of SSWS+ HPCSW*.

Perry, M. (2008). *A framework to support spatial, temporal and thematic analytics over semantic web data*. Retrieved from https://corescholar.libraries.wright.edu/etd_all/247/

Perry, M., Sheth, A., Arpinar, I. B., & Hakimpour, F. (2009). Geospatial and temporal semantic analytics. *Handbook of Research on Geoinformatics*, 161--170. Retrieved from https://books.google.ca/books?hl=en&lr=&id=CQ5If9TQQ-EC&oi=fnd&pg=PT196&dq=%22semantic+analytics+is%22&ots=aTPy2fzgJH&sig=sg5keTmlrFX_BBZp2m1oGNiqaDM

Pólya, G. (1954). *Mathematics and Plausible Reasoning (vol. 2): Patterns of plausible inference*. Princeton University Press.

Prade, H. (1985). A computational approach to approximate and plausible reasoning with applications to expert systems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *3*, 260–283. Retrieved from https://ieeexplore.ieee.org/abstract/document/4767656/

Predoiu, L., & Stuckenschmidt, H. (2008). Probabilistic models for the semantic web: A survey. *The Semantic Web for Knowledge and Data Management: Technologies and Practices*. Retrieved from https://books.google.ca/books?hl=en&lr=&id=4ot54J8g8eEC&oi=fnd&pg=PT90&dq=%22Semantic+Web+languages+are+based+on%22&ots=21M0yeD9Bw&sig=kF_XEXQvQfjzz31UbB-ZUl34cAk

Quach, X. H., & Hoang, T. L. G. (2018). Fuzzy Ontology Modeling by Utilizing Fuzzy Set and Fuzzy Description Logic. In *Modern Approaches for Intelligent Inf. and Database Systems*. Springer, Cham. http://doi.org/10.1007/978-3-319-76081-0_2

Ramakrishnan, C., Kochut, K. J., & Sheth, A. P. (2006). A framework for schema-driven relationship discovery from unstructured text. In *International Semantic Web Conference* (pp. 583--596). Retrieved from https://link.springer.com/10.1007%2F11926078_42

Rattanasawad, T., Buranarach, M., Saikaew, K. R., & Supnithi, T. (2018). A Comparative Study of Rule-Based Inference Engines for the Semantic Web. *IEICE TRANSACTIONS on Information and Systems*, *101*(1), 82--89. Retrieved from https://search.ieice.org/bin/summary.php?id=e101-d_1_82

Redavid, D., Iannone, L., Payne, T., & Semeraro, G. (2008). OWL-S Atomic Services Composition with SWRL Rules. In *Foundations of Intelligent Systems* (pp. 605–611). Berlin, Heidelberg: Springer Berlin Heidelberg. http://doi.org/10.1007/978-3-540-68123-6_66

Rindflesch, T., Kilicoglu, H., & Fiszman, M. (2011). Semantic MEDLINE: An advanced information management application for biomedicine. *Information Services & Use*, *31*(1–2), 15–21. Retrieved from http://content.iospress.com/articles/information-services-and-use/isu627

Rodríguez-González, A., Labra-Gayo, J. E., Colomo-Palacios, R., Mayer, M. A., Gómez-Berbís, J. M., & García-Crespo, A. (2012). SeDeLo: Using semantics and description logics to support aided clinical diagnosis. *Journal of Medical Systems*, *36*(4), 2471–2481. http://doi.org/10.1007/s10916-011-9714-1

Rosati, R. (2012). Prexto: Query rewriting under extensional constraints in DL− lite. *The Semantic Web: Research and Applications*.

Rosati, R., & Almatelli, A. (2010). Improving Query Answering over DL-Lite Ontologies. *KR*.

Roski, J., Bo-Linn, G. W., & Andrews, T. A. (2014). Creating Value In Health Care Through Big Data: Opportunities And Policy Implications. *Health Affairs*, *33*(7), 1115–1122. http://doi.org/10.1377/hlthaff.2014.0147

Sabou, M. (2016). An Introduction to Semantic Web Technologies. In *Semantic Web Technologies for Intelligent Engineering Applications* (pp. 53–81). Cham: Springer International Publishing. http://doi.org/10.1007/978-3-319-41490-4_3

Sadeghi, A., Lange, C., Vidal, M.-E., & Auer, S. (2017). Integration of Scholarly Communication Metadata Using Knowledge Graphs. In *International Conference on Theory and Practice of Digital Libraries* (pp. 328–341). Springer, Cham. http://doi.org/10.1007/978-3-319-67008-9_26

Schechter, L. M. (2015). A logic of plausible justifications. *Theoretical Computer Science*, *603*, 132--145. Retrieved from https://www.sciencedirect.com/science/article/pii/S0304397515006398

Schockaert, S., & Derrac, J. (2015). Commonsense Reasoning Based on Betweenness and Direction in Distributional Models. *2015 AAAI Spring Symposium Series*.

Schriml, L. M., Arze, C., Nadendla, S., Chang, Y.-W. W., Mazaitis, M., Felix, V., … Kibbe, W. A. (2012). Disease Ontology: a backbone for disease semantic integration. *Nucleic Acids Research*, *40*(D1), D940–D946. http://doi.org/10.1093/nar/gkr972

Serrano, M., & Gyrard, A. (2016). A Review of Tools for IoT Semantics and Data Streaming Analytics. *Building Blocks for IoT Analytics*, 139. Retrieved from https://books.google.ca/books?hl=en&lr=&id=uYktDwAAQBAJ&oi=fnd&pg=PA139&dq=%22semantic+analytics+is%22&ots=aflc7G5deH&sig=_33Glk4dVQy_KhvrtRamS4vjGoo

Shi, B., & Weninger, T. (2016). Discriminative Predicate Path Mining for Fact Checking in Knowledge Graphs. *Knowledge-Based Systems*, *104*, 123–133. http://doi.org/10.1016/j.knosys.2016.04.015

Shi, L., Li, S., Yang, X., Qi, J., Pan, G., & Zhou, B. (2017). Semantic health knowledge graph: Semantic integration of heterogeneous medical knowledge and services. *BioMed Research International*, 1--12. Retrieved from https://www.hindawi.com/journals/bmri/2017/2858423/abs/

Sion, A. (2009). A fortiori reasoning in judaic logic.

Ślęzak, D., Janusz, A., Świeboda, W., Nguyen, H. S., Bazan, J. G., & Skowron, A. (2011). Semantic analytics of PubMed content. In *Symposium of the Austrian HCI and Usability Engineering Group* (pp. 63--74). Retrieved from https://link.springer.com/10.1007%2F978-3-642-25364-5_7

Song, S. (2011). Metaphors and metonymy-A tentative research into modern cognitive linguistics. *Theory and Practice in Language Studies*, *1*(1), 68--73. Retrieved from http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.348.4624&rep=rep1&type=pdf#page=70

Stoilos, G., Stamou, G. B., Tzouvaras, V., Pan, J. Z., & Horrocks, I. (2005). Fuzzy OWL: Uncertainty and the Semantic Web. In *In: International Workshop of OWL: Experiences and Directions (OWL-ED 2005)*.

Tao, C., Zhang, Y., Jiang, G., Bouamrane, M.-M., & Chute, C. G. (2012). Optimizing semantic MEDLINE for translational science studies using semantic web technologies. In *Proceedings of MIXHS '12*. New York, New York, USA: ACM Press. http://doi.org/10.1145/2389672.2389683

Tickoo, O., & Iyer, R. (2017). Knowledge and Ontologies. In *Making Sense of Sensors* (pp. 83–94). Berkeley, CA: Apress. http://doi.org/10.1007/978-1-4302-6593-1_6

Tindale, C. (2010). *Reason's Dark Champions: Constructive Strategies of Sophistical Argument*. University of South Carolina Press.

Tsatsaronis, G., Balikas, G., Malakasiotis, P., & et. al. (2015). An overview of the BIOASQ large-scale biomedical semantic indexing and question answering competition. *BMC Bioinformatics*, *16*(1). http://doi.org/10.1186/s12859-015-0564-6

Tukey, J. W. (1977). *Exploratory data analysis*. Reading, Mass. Retrieved from https://pdfs.semanticscholar.org/2f40/fd06bc9fd00a27437e14ed171e96e4fd9326.pdf

Urbani, J., Van Harmelen, F., Schlobach, S., & Bal, H. (2011). QueryPIE: Backward reasoning for OWL Horst over very large knowledge bases. In *The Semantic Web--ISWC 2011*. Retrieved from https://link.springer.com/content/pdf/10.1007/978-3-642-25073-6_46.pdf

Van Ossenbruggen, J., Nack, F., & Hardman, L. (2004). That obscure object of desire: multimedia metadata on the web, part-1. *IEEE MultiMedia*, *11*(4), 38--48.

Van Woensel, W., Roy, P. C., & Abidi, S. S. R. (2016). SmartRL: A Context-Sensitive, Ontology-Based Rule Language for Assisted Living in Smart Environments. In *International Symposium on Rules and Rule Markup Languages for the Semantic Web* (pp. 341–349). Springer, Cham. http://doi.org/10.1007/978-3-319-42019-6_22

Vega Renon, L. (1998). Aristotle's Endoxa and Plausible Argumentation. *Argumentation*, *12*(1), 95–113. http://doi.org/10.1023/A:1007720902559

Velampalli, S., & Jonnalagedda, M. V. (2017). Graph based knowledge discovery using MapReduce and SUBDUE algorithm. *Data & Knowledge Engineering*, *111*, 103–113. http://doi.org/10.1016/J.DATAK.2017.08.001

Venetis, T., Stoilos, G., & Stamou, G. (2014). Query Extensions and Incremental Query Rewriting for OWL 2 QL Ontologies. *Journal on Data Semantics*, *3*(1), 1–23. http://doi.org/10.1007/s13740-012-0017-6

Virvou, M., & Kabassi, K. (2004). Adapting the human plausible reasoning theory to a graphical user interface. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, *34*(4), 546–563.

W3C. (n.d.). W3C Semantic Web Activity. Retrieved July 23, 2018, from https://www.w3.org/standards/semanticweb/query

Walton, D., Tindale, C. W., & Gordon, T. F. (2014). Applying Recent Argumentation Methods to Some Ancient Examples of Plausible Reasoning. *Argumentation*, *28*(1), 85–119. http://doi.org/10.1007/s10503-013-9306-y

Wang, C., Xu, P., Zhang, L., Huang, J., Zhu, K., & Luo, C. (2018). Current Strategies and Applications for Precision Drug Design. *Frontiers in Pharmacology*, *9*, 787. http://doi.org/10.3389/fphar.2018.00787

Wei, Y., Luo, J., & Xie, H. (2016). KGRL: An OWL2 RL Reasoning System for Large Scale Knowledge Graph. In *2016 12th International Conference on Semantics, Knowledge and Grids (SKG)* (pp. 83–89). IEEE. http://doi.org/10.1109/SKG.2016.020

Weil, A. R. (2014). Big Data In Health: A New Era For Research And Patient Care. *Health Affairs*, *33*(7), 1110–1110. http://doi.org/10.1377/hlthaff.2014.0689

Wishart, D. S., Feunang, Y. D., Guo, A. C., Lo, E. J., Marcu, A., Grant, J. R., … Wilson, M. (2018). DrugBank 5.0: a major update to the DrugBank database for 2018. *Nucleic Acids Research*, *46*(D1), D1074–D1082. http://doi.org/10.1093/nar/gkx1037

Woensel, W. Van, Mohammadhassanzadeh, H., Abidi, S., & Abidi, S. S. R. (2015). Multi-Strategy Semantic Web Reasoning for Medical Knowledge Bases. In *BDM2I*.

World Wide Web Consortium. (2012). OWL 2 web ontology language document overview.

Wu, Z., Eadon, G., & Das, S. (2008). Implementing an inference engine for RDFS/OWL constructs and user-defined rules in Oracle. *Data Engineering, IEEE 24th International Conference On*, 1239–1248.

Yang, Y., Huang, Z., Han, Y., Hua, X., & Tang, W. (2017). Using Knowledge Graph for Analysis of Neglected Influencing Factors of Statin-Induced Myopathy. In *International Conference on Brain Informatics* (pp. 304–311). Springer, Cham. http://doi.org/10.1007/978-3-319-70772-3_29

Yusof, M. M., Kuljis, J., Papazafeiropoulou, A., & Stergioulas, L. K. (2008). An evaluation framework for Health Information Systems: human, organization and technology-fit factors (HOT-fit). *International Journal of Medical Informatics*, *77*(6), 386–398. http://doi.org/10.1016/J.IJMEDINF.2007.08.011

Zhang, J., Yang, J., & Li, J. (2017). When Rule Engine Meets Big Data: Design and Implementation of a Distributed Rule Engine Using Spark. In *Big Data Computing Service and Applications (BigDataService), 2017 IEEE Third International Conference on* (pp. 41--49). Retrieved from https://ieeexplore.ieee.org/abstract/document/7944919/

Zhang, Y., Dai, H., Kozareva, Z., Smola, A. J., & Song, L. (2017). Variational Reasoning for Question Answering with Knowledge Graph. *ArXiv Preprint ArXiv:1709.04071*. Retrieved from http://arxiv.org/abs/1709.04071

Zhao, L., Munne, R. F., Kertkeidkachorn, N., & Ichise, R. (2017). Missing RDF Triples Detection and Correction in Knowledge Graphs. In *Joint International Semantic Technology Conference* (pp. 164–180). Springer, Cham. http://doi.org/10.1007/978-3-319-70682-5_11

Zhu, M., Qu, H., Zhao, J.-H., Chen, W., & Jalii, K. (2017). A methodology for quantitative analysis of non-determinacy relationships using fuzzy logic in Bayesian networks. In *2017 IEEE 2nd Information Technology, Networking, Electronic and Automation Control Conference (ITNEC)* (pp. 856–862). IEEE. http://doi.org/10.1109/ITNEC.2017.8284857

Zimmermann, A., Lopes, N., Polleres, A., & Straccia, U. (2012). A general framework for representing, reasoning and querying with annotated semantic web data. *Web Semantics: Science, Services and Agents on the World Wide Web*, *11*, 72--95. Retrieved from https://www.sciencedirect.com/science/article/pii/S1570826811000771

Zurek, T. (2012). Modelling of a’fortiori reasoning. *Expert Systems with Applications*, *39*(12), 10772–10779.

Appendix I: PL-OWL Extension

```
@prefix dc: <http://purl.org/dc/elements/1.1/> .
@prefix grddl: <http://www.w3.org/2003/g/data-view#> .
@prefix owl: <http://www.w3.org/2002/07/owl#> .
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
@prefix xml: <http://www.w3.org/XML/1998/namespace> .
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .
@prefix plowl: <http://niche.cs.dal.ca/2017/06/plowl#> .

<http://niche.cs.dal.ca/2017/06/plowl/> a owl:Ontology ;
     dc:title "Plausible OWL Schema vocabulary" ;
     rdfs:comment """
                 This ontology partially describes the plausible extensions to OWL. This extension
includes
                 classes and properties that together form the basis of the PL-OWL.
                 """ ;
     rdfs:isDefinedBy
          <http://www.w3.org/TR/owl2-mapping-to-rdf/>,
          <http://www.w3.org/TR/owl2-rdf-based-semantics/>,
          <http://www.w3.org/TR/owl2-syntax/> ;
     rdfs:seeAlso   <http://www.w3.org/TR/owl2-rdf-based-semantics/#table-axiomatic-classes>,
                    <http://www.w3.org/TR/owl2-rdf-based-semantics/#table-axiomatic-properties> ;
     owl:imports <http://www.w3.org/2000/01/rdf-schema>,
                                       <http://www.w3.org/2002/07/owl> ;
     owl:versionIRI <http://niche.cs.dal.ca/2017/06/plowl> ;
     owl:versionInfo "$Date: 2016/11/07 14:59:12 $" ;
     grddl:namespaceTransformation <http://dev.w3.org/cvsweb/2009/owl-grddl/owx2rdf.xsl> .

plowl:OrderedProperty
        a owl:Class ;
         rdfs:label "OrderedProperty" ;
         rdfs:comment "The class of ordered properties." ;
         rdfs:isDefinedBy <http://web.cs.dal.ca/~hossein/plowl/plowl#> ;
         rdfs:subClassOf owl:ObjectProperty.

plowl:Context
        a owl:Class;
        rdfs:subClassOf rdf:Node;
        rdfs:label "Context" ;
    rdfs:comment "Represents the context of the object property." ;
    rdfs:isDefinedBy <http://web.cs.dal.ca/~hossein/plowl/plowl#> .

plowl:standsAfter
        a plowl:OrderedProperty;
         rdfs:label "standsAfter" ;
         rdfs:comment "The property that represents the partial order of two classes." ;
         rdfs:isDefinedBy <http://web.cs.dal.ca/~hossein/plowl/plowl#> .

plowl:PlausiblePattern
        a owl:Class;
         rdfs:subClassOf rdf:Node;
         rdfs:label "Plausible Pattern" ;
         rdfs:comment "Collection of the plausible patterns." ;
        rdfs:isDefinedBy <http://web.cs.dal.ca/~hossein/plowl/plowl#> .

plowl:PlausibleAnswer
        a owl:Class;
         rdfs:label "Plausible Answer" ;
         rdfs:comment "Represent the plasubily inferred triples." ;
         rdfs:isDefinedBy <http://web.cs.dal.ca/~hossein/plowl/plowl#>;
         rdfs:subClassOf [
                 a owl:Restriction ;
                 owl:onProperty plowl:inferredThroughPattern ;
                 owl:allValuesFrom [
                         a owl:Class ;
                         owl:unionOf (
                                 <http://web.cs.dal.ca/~hossein/plowl/univ-
bench.owl#Generalization>
                                 <http://web.cs.dal.ca/~hossein/plowl/univ-
bench.owl#Specialization>
                                 <http://web.cs.dal.ca/~hossein/plowl/univ-bench.owl#Similarity>
                                 <http://web.cs.dal.ca/~hossein/plowl/univ-
bench.owl#Dissimilarity>
                                 <http://web.cs.dal.ca/~hossein/plowl/univ-bench.owl#Afortiori>
                                 <http://web.cs.dal.ca/~hossein/plowl/univ-
bench.owl#Interpolation>
                         )
                         ]
```

```
                 ],
                 [
                 a owl:Restriction ;
                 owl:onProperty plowl:inTheContextOf ;
                 owl:someValuesFrom plowl:Context
                 ];

plowl:hasContext
        a owl:ObjectProperty;
          rdfs:range plowl:Context;
          rdfs:comment "This object property links an opject property to the context nodes being
applied to it."^^xsd:string ;
          rdfs:isDefinedBy <http://web.cs.dal.ca/~hossein/plowl/plowl#> .

plowl:standsAfter
        a owl:OrderedProperty;
          rdfs:range owl:Thing;
          rdfs:domain owl:Thing;
          rdfs:comment "This object property is used to model ordering relation to show which
concept (subject) locates after another concept (object) regarding a specific context. The inverse
property is standsBefore.."^^xsd:string ;
          rdfs:isDefinedBy <http://web.cs.dal.ca/~hossein/plowl/plowl#> .

plowl:standsBefore
        a owl:OrderedProperty;
          rdfs:range owl:Thing;
          rdfs:domain owl:Thing;
          rdfs:comment "This object property is used to model ordering relation to show which
concept (subject) locates before another concept (object) regarding a specific context. The inverse
property is standsAfter.."^^xsd:string ;
          rdfs:isDefinedBy <http://web.cs.dal.ca/~hossein/plowl/plowl#> .

plowl:hasSubject
        a owl:ObjectProperty;
          rdfs:domain plowl:PlausibleAnswer;
          rdfs:range  owl:Thing;
          rdfs:comment "This object property links a plausible answer to its subject."^^xsd:string
;
          rdfs:isDefinedBy <http://web.cs.dal.ca/~hossein/plowl/plowl#> .

plowl:hasObject
        a owl:ObjectProperty;
          rdfs:domain plowl:PlausibleAnswer;
          rdfs:range  owl:Thing;
          rdfs:comment "This object property links a plausible answer to its object."^^xsd:string ;
          rdfs:isDefinedBy <http://web.cs.dal.ca/~hossein/plowl/plowl#> .

plowl:hasPredicate
        a owl:ObjectProperty;
          rdfs:domain plowl:PlausibleAnswer;
          rdfs:range  owl:Thing;
          rdfs:comment "This object property links a plausible answer to its
predicate."^^xsd:string ;
          rdfs:isDefinedBy <http://web.cs.dal.ca/~hossein/plowl/plowl#> .

plwol:inferredThroughPattern
        a owl:ObjectProperty;
          rdfs:domain plowl:PlausibleAnswer;
          rdfs:range  owl:PlausiblePattern;
          rdfs:comment "This object property links a plausible answer to the plausible pattern that
lead to the inference."^^xsd:string ;
          rdfs:isDefinedBy <http://web.cs.dal.ca/~hossein/plowl/plowl#> .

plwol:inTheContextOf
        a owl:ObjectProperty;
          rdfs:domain plowl:PlausibleAnswer;
          rdfs:range  plowl:Context;
          rdfs:comment "This object property links a plausible answer to the context that lead to
the inference."^^xsd:string ;
          rdfs:isDefinedBy <http://web.cs.dal.ca/~hossein/plowl/plowl#> .
```

Appendix II: List of initially answered questions (without query modification)

#	Question	Source	Domain	Type
1	What causes Scurvy?	Train	cause	Factoid
2	What is known as the cause of subacute thyroiditis?	Train	cause	Factoid
3	Which viruses are best known to cause myocarditis?	Train	cause	Factoid
4	Which bacteria caused plague?	Train	cause	Factoid
5	Which pituitary adenoma is common cause of infertility in women?	Train	cause	Factoid
6	Which could be some of the possible causes of hypersomnia?	Train	cause	Factoid
7	Is valproic acid effective for glioblastoma treatment?	Train	treatment	Yes/No
8	What is the treatment of acute pericarditis?	Train	treatment	Factoid
9	What is the treatment of subacute thyroiditis?	Train	treatment	Factoid
10	Which disease can be treated with Delamanid?	Train	treatment	Factoid
11	What are the treatments for GIST (gastrointestinal stromal tumor)?	Train	treatment	Factoid
12	What is the treatment of Riedel disease (thyroiditis)?	Train	treatment	Factoid
13	Which drugs are utilized to treat eosinophilic esophagitis?	Train	treatment	Factoid
14	What is the treatment of choice for gastric lymphoma?	Train	treatment	Factoid
15	What is the treatment of neuropathic pain in children?	Train	treatment	Factoid
16	Which acetylcholinesterase inhibitors are used for treatment of myasthenia gravis?	Train	treatment	Factoid
17	Which is the most common cause of sudden cardiac death in young athletes?	Train	cause	Factoid
18	Which deficiency is the cause of restless leg syndrome?	Train	cause	Factoid
19	Mutation of which gene and which chromosome cause Neurofibromatosis type I?	Train	cause	Factoid
20	What organism causes tularemia?	Train	cause	Factoid
21	What is the cause of Tardive dyskinesia?	Train	cause	Factoid
22	Which is the main cause of the Patau syndrome?	Train	cause	Factoid
23	Which bacteria cause diphtheria?	Train	cause	Factoid
24	How is primary intestinal lymphangiectasia (PIL) caused?	Train	cause	Factoid
25	What fruit causes Jamaican vomiting sickness?	Train	cause	Factoid
26	Which virus type causes Molluscum contagiosum?	Train	cause	Factoid
27	Which are the main brain dysfunctions caused by hyperbilirubinemia?	Train	cause	Factoid
28	Which are the causes of the Koebner phenomenon?	Train	cause	Factoid
29	Which gene mutations cause the Marfan syndrome?	Train	cause	Factoid
30	Is propranolol used for treatment of infantile hemangioma?	Testset1	treatment	Yes/No
31	Is enzastaurin effective treatment of glioblastoma?	Testset1	treatment	Yes/No
32	Mutations in which gene cause Schimke immune-osseous dysplasia?	Testset2	cause	Factoid
33	Which disorder has been approved for treatment with Alk inhibitors?	Testset2	treatment	Factoid
34	Can radius fracture cause carpal tunnel syndrome?	Testset3	cause	Yes/No
35	Is cilengitide effective for treatment of glioblastoma?	Testset3	treatment	Yes/No
36	What is the first line treatment for sarcoidosis?	Testset3	treatment	Factoid
37	List 2 approved drug treatments for Inflammatory Bowel Disease (IBD).	Testset3	treatment	Factoid
38	Which disease is treated with Fexinidazole?	Testset4	treatment	Factoid
39	What organism causes scarlet fever also known as scarletina?	Testset4	cause	Factoid
40	Is subacute sclerosing panencephalitis caused by the Measles vaccine?	Testset4	cause	Yes/No
41	What drug treatment can cause a spinal epidural hematoma?	Testset5	cause	Factoid
42	List diseases caused by protein glutamine expanded repeats.	Testset5	cause	Factoid
43	What causes leishmaniasis?	Testset5	cause	Factoid
44	Please list 10 conditions which play a role in causing atrial fibrillation.	Testset5	cause	Factoid
45	What drug treatment can cause a spinal epidural hematoma?	Testset5	treatment	Factoid
46	Please list 3 diseases treated with Valtrex(valacyclovir)	Testset5	treatment	Factoid
47	Which enzyme deficiency can cause GM1 gangliosidoses?	Train	cause	Factoid
48	What is the treatment of triiodothyronine toxicosis (T3_thyrotoxicosis)?	Train	treatment	Factoid
49	List FDA approved treatments for androgenetic allopecia.	Train	treatment	Factoid
50	List 4 drugs used to treat opioid addiction or overdose.	Testset4	treatment	Factoid
51	Which drugs are utilized to treat eosinophilic esophagitis?	Train	treatment	Factoid
52	Which drugs have been found effective for the treatment of chordoma?	Train	treatment	Factoid

Appendix III: List of plausibly answered questions

#	Query	Source	Domain	Type
1	What causes Katayama Fever?	Train	cause	Factoid
2	What is the cause of episodic ataxia type 6?	Train	cause	Factoid
3	Do statins cause diabetes?	Train	cause	Yes/No
4	Can Levoxyl (levothyroxine sodium) cause insomnia?	Train	cause	Yes/No
5	Which antibodies cause Riedel thyroiditis?	Train	cause	Factoid
6	Is the monoclonal antibody Trastuzumab (Herceptin) of potential use in the treatment of prostate cancer?	Train	treatment	Yes/No
7	What is the treatment of acute myocarditis?	Train	treatment	Factoid
8	What is the genetic basis of the Delayed Sleep-Phase Syndrome (DSPS)?	Testset1	cause	Factoid
9	Does DDX54 play a role in DNA damage response?	Testset1	cause	Yes/No
10	Does a tonsillectomy affect the patient's voice?	Testset1	cause	Yes/No
11	Is there an RNAi drug being developed to treat amyloidosis?	Testset1	treatment	Yes/No
12	Are there RNAi approaches considered for the treatment of kidney injury?	Testset1	treatment	Yes/No
13	Has IVIG been tested in clinical trials for the treatment of Alzheimer's disease?	Testset1	treatment	Yes/No
14	Which bacteria causes erythrasma?	Testset2	cause	Factoid
15	Do bacteria from the genus Morexella cause respiratory infections?	Testset2	cause	Yes/No
16	Was saracatinib being considered as a treatment for Alzheimer's disease in November 2017?	Testset2	treatment	Yes/No
17	Is celiac disease caused by gliadin-induced transglutaminase-2 (TG2)-dependent events?	Testset3	cause	Yes/No
18	Can doxycycline cause photosensitivity?	Testset3	cause	Yes/No
19	What causes Black Lung?	Testset3	cause	Factoid
20	Can canagliflozin cause euglycemic diabetic ketoacidosis?	Testset3	cause	Yes/No
21	Mutation of which gene causes arterial tortuosity syndrome?	Testset3	cause	Factoid
22	Can CD55 deficiency cause thrombosis?	Testset4	cause	Yes/No
23	Which diseases are caused by mutations in Calsequestrin 2 (CASQ2) gene?	Train	cause	Factoid
24	List disorders that are caused by mutations in the mitochondrial MTND6 gene.	Train	cause	Factoid
25	What organism causes woolsorter's disease	Train	cause	Factoid
26	Which disease(s) are caused by HEX A deficiency?	Train	cause	Factoid
27	Is Brucella abortus the organism that causes brucillosis known to cause spontaneous abortions in humans?	Testset5	cause	Yes/No
28	Are AAV vectors considered for the treatment of retinal dystrophies?	Testset5	treatment	Yes/No
29	Dinutuximab is used for treatment of which disease?	Testset5	treatment	Factoid
30	What is the cause of Phthiriasis Palpebrarum?	Train	cause	Factoid
31	Orteronel was developed for treatment of which cancer?	Train	treatment	Factoid
32	Matuzumab has been tested for treatment of which cancers?	Train	treatment	Factoid
33	Is nivolumab used for treatment of Non-Small-Cell Lung Cancer?	Train	treatment	Yes/No
34	Is lambrolizumab effective for treatment of patients with melanoma?	Train	treatment	Yes/No
35	Which diseases can be treated with Afamelanotide?	Train	treatment	Factoid
36	List the diseases that can be treated using Vedolizumab.	Train	treatment	Factoid
37	Is Migalastat used for treatment of Fabry Disease?	Train	treatment	Yes/No
38	Is ocrelizumab effective for treatment of multiple sclerosis?	Train	treatment	Yes/No
39	For the treatment of which conditions can atypical neuroleptic drugs be used?	Train	treatment	Factoid
40	Is tretinoin effective for photoaging?	Testset1	treatment	Yes/No
41	Could Arimidex (anastrozole) cause hot flashes? (hot flushes)	Train	cause	Yes/No
42	What is the definitive treatment for low pressure headache?	Train	treatment	Factoid

Appendix IV: List of remained unanswered questions

#	Query	Source	Domain	Type
1	Which are the main causes of fetal echogenic bowel?	Train	cause	Factoid
2	Can vitamin B1 deficiency cause encephalopathy?	Train	cause	Yes/No
3	What causes erucism?	Train	cause	Factoid
4	Which gene-defect causes the Vel-blood type?	Train	cause	Factoid
5	What is the treatment of amiodarone-induced thyrotoxicosis?	Train	treatment	Factoid
6	List all reported treatment options for anxiety in autism spectrum disorder.	Train	treatment	Factoid
7	Dracorhodin perchlorate was tested for treatment of which cancers?	Train	treatment	Factoid
8	Is armodafinil used for treatment of insomnia?	Train	treatment	Yes/No
9	Pridopidine has been tested for treatment of which disorder?	Train	treatment	Factoid
10	Which disease is treated with Eliglustat?	Train	treatment	Factoid
11	What is the treatment of interferon-induced thyroiditis?	Train	treatment	Factoid
12	Which inherited disorder is known to be caused by mutations in the NEMO gene?	Train	cause	Factoid
13	What causes Puffy hand syndrome?	Testset2	cause	Factoid
14	What protein is the most common cause of hereditary renal amyloidosis?	Testset2	cause	Factoid
15	Centor criteria are used for which disease?	Testset2	treatment	Factoid
16	Which personality disorder is treated using dialectical behavior therapy?	Testset3	treatment	Factoid
17	Milwaukee protocol was tested for treatment of which disease?	Testset4	treatment	Factoid
18	Does SARM1 deletion cause neurodegeneration?	Testset4	cause	Yes/No
19	A bite from the Lone Star Tick Amblyomma americanum, can cause the victim to become allergic to red meat, yes or no?	Testset4	cause	Yes/No
20	What is caused by the ectopic expression of CTCF?	Testset5	cause	Factoid

Appendix V: Average of execution time of SPARQL queries and rewriting time of a plausible query time by each question

#	SPARQL query execution time		Rewriting time	
	Average	StdDev	Average	StdDev
1	3724	9615.7	15	11.8
2	4819	213261.4	15	12.7
3	3708	7498.5	17	27.9
4	4743	76696.3	9	53.8
5	4641	13167.5	13	9.9
6.1	3605	4936.5	13	11.5
6.2	3358	5681.0	19	29.8
7	5135	5304.9	12	11.8
8	3928	9122.5	19	10.6
9.2	4404	9849.1	10	61.6
10	4291	10044.3	10	20.4
11	5148	12305.9	9	46.3
12	4945	11536.2	9	71.3
13	5607	8147.2	11	26.3
14	5240	28602.0	17	14.8
15.1	4837	5912.6	12	19.3
15.2	3104	5083.2	33	103.7
16.1	5326	7721.0	11	33.1
17.1	2969	6387.1	21	14.1
17.2	2809	6677.3	16	11.4
18.1	3310	5948.1	11	16.7
18.2	3902	12100.3	14	24.0
19	9537	19063.9	14	11.0
20.1	3780	5641.1	13	25.9
20.2	5126	12157.1	14	14.2
21	6650	9574.0	11	11.7
22	3062	5188.2	20	62.5
23	5895	8279.8	17	11.8
24	11843	48224.9	13	6.3
25	3409	5911.1	15	12.2
26	3920	6034.8	12	8.1
27.1	4092	4496.0	12	20.7
27.2	2806	3466.5	22	52.4
28.1	3844	3936.8	12	18.1
28.2	3369	5839.1	15	25.4
29	12819	32956.3	18	10.2
30	3918	5758.0	14	30.3
31	4214	10506.0	16	8.1
32	9369	21085.8	23	14.0
33	3637	5760.8	41	141.6
34	3261	3523.5	23	13.2
35	31156	83894.4	19	10.7
36	7288	15861.5	24	12.3
37.1	3159	6550.5	19	9.9
37.2	4415	5400.6	12	33.7
38.1	18715	58897.1	226	420.7
38.2	3721	6802.6	42	206.9
39	3997	5684.0	15	9.7
40.1	4407	9127.2	13	10.9
41.1	2892	5514.6	19	15.1
42	5618	11356.9	12	14.9

Appendix VI: Retrieval time of applicable plausible semantics by plausible patterns

No.	AFORT (More to Less)		AFORT (Less to More)		GEN		SPEC		SIM (Hierarchical)		SIM (Ordered)		DIS (Hierarchical)		DIS (Ordered)		INTP (Stands Before)		INTP (Stands After)	
	Avg.	StdDev	Avg.	StdDev	Avg.	StdDev	Avg.	StdDev	Avg.	StdDev	Avg.	StdDev	Avg.	StdDev	Avg.	StdDev	Avg.	StdDev	Avg.	StdDev
1	384.4	1302.4	12.5	15.1	13.0	10.2	7.6	8.6	13.8	20.8	27.0	33.3	8.8	9.5	4.4	1.7	35.7	91.8	4.5	1.5
2	980.1	6757.6	8.4	6.1	9.5	5.3	9.1	20.7	9.1	8.0	16.9	28.3	7.2	7.2	4.9	1.7	16.5	51.1	4.6	1.6
3	122.5	124.9	5.7	3.4	9.3	19.3	17.7	38.3	6.1	5.7	18.8	28.5	6.1	4.6	5.6	6.4	15.3	76.3	4.9	1.8
4	53.7	208.9	46.0	212.0	39.0	186.2	43.1	215	41.2	185	44.0	209	47.3	210	45	218	34.0	181.1	36.6	188
5	128.1	166.7	10.2	10.4	11.4	10.3	4.3	1.8	11.8	11.8	10.2	9.0	6.5	7.0	3.8	1.0	4.9	3.7	4.2	2.3
6.1	95.7	64.0	5.6	2.3	11.1	16.3	6.6	6.3	5.2	1.6	20.7	29.9	5.2	1.4	5.3	2.0	22.8	60.3	5.6	2.5
6.2	133.4	154.0	5.4	3.9	21.0	59.0	18.6	31.8	5.9	5.4	31.6	48.4	5.5	2.7	5.8	6.4	27.5	111.0	4.9	2.9
7	55.5	70.1	7.5	3.1	11.5	18.9	7.5	4.5	7.5	12.5	15.8	20.7	6.8	2.8	6.4	3.5	17.4	56.0	7.9	3.4
8	174.1	245.7	15.1	10.7	40.0	103.2	6.2	4.6	14.7	10.9	17.9	34.4	10.9	14.9	6.1	3.6	12.5	27.8	6.1	4.6
9.1	71.3	89.3	9.7	7.1	14.1	19.9	8.8	4.9	9.5	13.7	18.5	29.1	8.1	3.3	6.0	4.1	18.9	68.4	9.2	4.3
9.2	27.9	53.7	3.0	3.8	3.8	9.8	3.7	9.1	3.0	3.4	4.8	12.4	3.0	4.4	2.9	2.6	3.3	15.2	3.0	4.0
10	85.3	87.1	5.3	5.6	7.2	11.5	6.2	8.8	5.2	7.8	11.5	19.0	5.0	4.5	4.8	2.8	6.7	20.4	4.9	3.0
11	32.4	130.0	4.9	12.9	5.4	15.0	4.7	8.5	4.6	11.6	6.7	15.1	4.6	13.3	4.5	8.6	4.9	10.8	4.7	11.3
12	29.9	72.1	4.5	7.8	5.2	15.2	4.5	8.1	4.4	9.7	6.2	11.9	4.2	8.3	4.3	6.8	4.9	12.5	4.4	8.5
13	69.0	70.2	9.9	11.4	13.1	22.0	9.3	8.4	9.2	7.1	15.8	23.2	7.9	3.4	6.9	6.2	13.5	51.4	9.1	4.5
14	2543.7	6560.7	36.0	30.0	35.3	23.8	14.7	16.4	82.4	146	34.8	29.8	30.2	47	5.0	2.1	14.1	25.7	4.9	2.4
15.1	67.4	52.2	7.4	3.8	12.3	21.8	7.7	9.6	7.7	12.3	14.3	16.9	6.9	6.1	6.2	3.4	15.6	62.4	7.4	3.4
15.2	524.3	1154.7	8.8	3.4	7.0	3.0	8.0	3.9	7.6	4.3	29.1	14.9	6.7	2.6	6.7	3.4	22.8	35.0	6.0	2.1
16.1	66.8	77.3	9.1	5.8	12.0	17.4	8.6	7.4	8.5	4.9	15.5	23.8	7.7	3.4	6.6	4.7	12.8	55.6	8.5	4.1
17.1	168.2	0.0	25.1	0.0	3.9	0.0	4.1	0.0	8.8	0.0	28.8	0.0	26.5	0.0	3.9	0.0	4.1	0.0	3.6	0.0
17.2	194.8	106.1	8.9	9.1	28.1	52.5	19.0	37.6	5.9	7.4	44.1	62.7	7.3	9.2	3.2	1.2	9.1	11.3	4.5	3.9
18.1	87.6	61.2	5.6	3.9	11.4	18.4	5.3	5.9	8.5	19.3	13.1	17.7	4.5	2.0	4.5	2.3	20.7	129.0	4.2	1.5
18.2	129.6	212.9	6.3	5.7	10.6	23.8	9.6	21.8	6.6	13.8	16.5	23.4	5.4	2.6	5.0	1.7	10.8	38.8	5.0	2.0
19	44.5	60.0	5.3	1.3	11.6	13.1	6.2	1.5	9.9	9.6	6.4	1.3	5.5	0.9	6.0	1.8	6.6	1.4	5.5	2.0
20.1	51.9	67.5	4.9	2.1	6.4	11.1	6.0	8.1	5.2	3.3	10.4	15.1	5.4	6.0	4.9	1.7	7.4	28.1	5.0	2.4
20.2	55.7	69.9	7.5	1.8	8.2	1.4	5.9	0.3	4.6	2.7	5.7	0.8	4.7	2.5	6.9	0.2	6.7	1.6	10.3	8.8
21	38.1	16.8	5.4	0.8	9.8	7.4	5.3	2.0	6.1	0.6	9.1	5.9	5.0	1.5	5.5	2.3	12.3	11.5	6.5	1.2
22	87.1	175.5	5.1	4.7	17.4	24.9	7.2	7.8	4.4	2.4	16.3	22.4	4.6	3.2	4.5	4.3	6.0	6.4	4.6	4.4
23	68.1	60.8	3.3	2.6	3.8	3.0	2.5	1.9	4.3	3.7	5.5	6.0	2.9	2.3	3.3	2.7	5.1	5.4	2.9	2.3
24	194.9	0.0	6.8	0.0	9.2	0.0	3.7	0.0	4.3	0.0	5.6	0.0	4.0	0.0	4.0	0.0	4.0	0.0	4.3	0.0

25	92.8	75.9	7.2	7.4	12.8	10.9	6.7	10.1	14.5	31.1	14.5	15.3	3.9	0.7	4.0	2.0	10.3	14.1	4.0	1.1
26	69.0	42.4	7.2	7.5	12.7	28.9	7.5	8.3	5.7	3.0	18.5	24.5	5.4	1.2	5.0	1.5	30.3	168.4	5.4	1.7
27.1	51.0	65.7	6.3	3.0	8.9	13.1	6.5	5.1	6.3	3.2	12.7	14.2	5.6	1.7	6.0	2.9	13.3	34.4	6.3	2.6
27.2	74.6	235.4	5.9	5.2	6.6	7.9	8.8	10.5	5.8	3.3	19.8	20.0	4.9	1.9	5.2	3.1	20.2	40.6	4.6	1.6
28.1	25.5	36.1	6.7	8.4	6.3	2.8	7.3	7.6	6.2	2.5	13.5	14.1	5.9	2.5	5.7	3.3	14.9	51.1	6.5	3.7
28.2	78.5	93.1	4.6	2.4	6.9	13.8	9.5	24.6	5.5	8.1	11.9	19.4	4.7	2.3	4.4	2.0	8.7	37.6	4.3	1.7
29	31.9	76.6	5.0	1.3	6.2	5.5	21.8	55.2	15.3	33.0	13.7	28.4	4.6	0.6	4.2	0.7	6.0	3.2	5.6	4.4
30	66.4	147.0	6.2	2.9	7.4	6.5	6.8	6.6	5.3	2.1	18.3	27.9	5.9	2.3	5.8	1.6	52.6	209.7	5.6	1.8
31	114.0	96.2	5.1	2.6	7.5	5.5	4.4	1.2	5.6	3.8	12.6	17.9	4.4	1.0	4.4	1.0	6.0	7.9	4.2	1.0
32	49.1	97.4	9.9	0.6	8.5	2.0	11.1	7.5	7.4	1.7	13.6	12.9	9.5	4.9	7.7	2.5	10.3	4.8	12.2	11.4
33	15.9	22.8	8.4	2.2	35.6	74.4	27.6	53.9	7.5	1.5	29.4	56.2	6.7	1.6	6.5	1.6	9.4	5.6	6.7	1.6
34	5.2	1.2	5.2	1.5	5.3	1.5	7.5	8.3	5.3	1.2	9.6	13.3	6.7	4.9	7.3	7.6	7.0	6.3	4.7	0.5
35	447.0	881.3	14.7	17.8	5.1	0.8	11.0	11.4	6.1	0.3	38.9	67.7	5.5	0.9	5.0	0.9	5.7	0.6	5.5	0.6
36	7.0	1.1	8.1	1.7	6.9	2.1	8.1	4.1	7.0	0.7	9.6	8.2	6.8	1.2	6.6	0.9	8.0	5.5	6.6	1.7
37.1	54.9	69.9	6.8	7.1	5.6	5.3	5.9	5.4	4.4	1.0	10.9	8.9	4.1	0.6	4.0	0.8	4.7	1.4	4.2	0.9
37.2	58.1	77.2	7.4	3.6	9.8	8.8	7.2	5.0	6.9	3.3	14.4	15.7	6.4	2.8	7.0	3.9	23.5	100.3	7.0	3.3
38.1	14.5	10.1	9.7	2.1	9.1	4.9	9.6	4.2	13.6	4.0	10.8	1.0	8.1	0.5	10.2	1.8	7.7	0.5	168.1	223.5
38.2	106.7	301.0	12.0	2.1	14.2	9.7	15.0	18.4	10.4	2.0	15.3	15.5	9.7	2.0	9.1	2.3	16.0	15.5	9.2	2.4
39	86.6	57.4	5.0	2.3	11.6	30.3	9.1	20.0	5.6	3.9	16.5	32.6	4.8	1.4	4.7	1.6	16.2	67.0	4.7	1.7
40.1	22.9	27.0	5.5	1.3	7.4	7.6	4.9	0.6	4.1	1.2	6.0	3.9	4.4	1.0	13.6	15.9	4.7	0.5	4.2	0.8
41.1	209.2	108.5	5.9	3.0	10.1	11.9	16.3	35.9	5.5	1.3	38.2	84.7	4.6	0.8	4.3	0.7	59.8	184.8	4.3	0.7
42	81.8	94.6	9.4	6.4	15.8	32.7	10.4	24.4	8.7	6.7	23.0	54.6	7.4	4.6	6.9	4.1	17.2	60.4	8.7	4.0

Appendix V: Notices of permission to use excerpts from author’s publications

In this thesis, large and small excerpts were taken from three of the author's own published papers. Two of the manuscripts (Mohammadhassanzadeh, Raza Abidi, et al., 2017) and (Mohammadhassanzadeh, Van Woensel, et al., 2017) have been published in Biodata mining journal and CEUR Workshop Proceedings, respectively, which are free open-access publication services. The third manuscripts (Mohammadhassanzadeh et al., 2018) has been published in an IEEE proceeding, which does not require individuals working on their own thesis to obtain a formal reuse license (https://www.ieee.org/content/dam/ieee-org/ieee/web/org/pubs/permissions_faq.pdf).

Also, there are parts from another paper which is in progress at the time of submitting this thesis and will be submitted soon. A form of the student’s contribution to the manuscript was signed and submitted to the Faculty of Graduate Studies.