Information Retrieval with Actions and Change: an ASP-Based Solution

MARCELLO BALDUCCINI, EMILY C. LEBLANC
Drexel University, Philadelphia, PA

submitted 1 January 2003; revised 1 January 2003; accepted 1 January 2003

Abstract

Information Retrieval (IR) aims at retrieving documents that are most relevant to a query provided by a user. Traditional techniques rely mostly on syntactic methods. In some cases, however, links at a deeper semantic level must be considered. In this paper, we explore a type of IR task in which documents describe sequences of events, and queries are about the state of the world after such events. In this context, successfully matching documents and query requires considering the events' possibly implicit, uncertain effects and side-effects. We begin by analyzing the problem, then propose an action language based formalization, and finally automate the corresponding IR task using Answer Set Programming.

KEYWORDS: Reasoning about Actions and Change, Answer Set Programming, Information Retrieval

1 Introduction

Information Retrieval (IR) (Korfhage 1997) aims at identifying, among a set of available information sources, those that are most relevant to a query provided by the user. IR is arguably a staple of every day life – we consult Wikipedia for general reference, doctors search private databases for patient information, and researchers use public databases to find scientific publications. IR is also at the core of numerous commercial activities such as searching for news about business partners or competitors.

Most IR systems base the relevance of a source on a syntactic measurement of the overlap of terms between query and source (Manning et al. 2008). Even advanced techniques still focus on syntactic matching, and include temporal ordering (Campos 2015), query expansion (Carpineto and Ramano 2012), and graph based term weighting (Blanco and Lioma 2012).

However, research has demonstrated (Glavas and Snajder 2014) that traditional IR yields low accuracy when applied to documents centered on events, such as police reports, medical records, and breaking news. As one can imagine, documents of these kinds occur in large quantities and often contain very valuable information.

(Glavas and Snajder 2014) proposed a new approach, called event-centered IR, which succeeded in increasing match accuracy by means of some level of semantic analysis. However, their approach was limited to matching events mentioned in both queries and sources.

In this paper, we advance this line of research by considering the case in which the goal is to match sources containing sequences of events with queries that are about the *state of the world* after those events. This is the case, for example, in which the sources describe the history of a domain (e.g., historical documents, police reports, computer event logs) and a user is looking for sources from which the state of the domain at a moment of interest can be reconstructed (e.g., "was the firewall on when the attack happened?"). Our approach aims to identify reasonable matches even when a definitive answer cannot be immediately found in the sources, events have complex/hidden effects, and information is incomplete.

We begin by analyzing the problem and, appealing to commonsense and intuition, determine reasonable outcomes of the task as a human reader might carry it out. We use toy examples, which we progressively elaborate, but the approach easily applies to practical cases. Later, we develop needed mathematical foundations and propose a formalization of the reasoning task. It should be noted that, in this paper, we assume that passages in natural language have already been translated into a suitable logic form. The natural language task is orthgonal to the problem addressed in this paper, and will be considered at a later stage. Let us start from the following:

Example 1
The user's query, $\mathcal{Q}$, is "Is John married?" Available information sources are:
$\mathcal{S}_1$*: "John went on his first date with Mary."*
$\mathcal{S}_2$*: "John read a book."*
We want to determine which source is most relevant to $\mathcal{Q}$*.*

The query refers to the current state of the world, which with some approximation we can identify with the final state of the world in the sources. The sources describe events that occurred over time. Neither source mentions being married, making syntactic-based methods unfit for the task. However, from an intuitive perspective, $\mathcal{S}_1$ is more relevant to $\mathcal{Q}$ than $\mathcal{S}_2$. In fact, $\mathcal{S}_1$, together with commonsense knowledge that married people (normally) do not go on first dates, provides a strong indication that John is not married. $\mathcal{S}_2$, on the other hand, provides no information pertaining the query.

In this simple example, one can not only identify $\mathcal{S}_1$ as the most relevant source, but also obtain an accurate answer to the question. The simplicity of the example blurs the line between IR and question answering. In general, however, providing an accurate answer requires a substantial amount of reasoning to be carried out once a relevant source has been identified, as well as deep understanding of the content of the source and a large amount of world knowledge – something that is still challenging for state-of-the-art approaches. Thus, in this paper, *we assume that a reader with human-level intelligence will later find accurate answers by studying the sources identified as relevant by our approach. We focus on defining techniques that provide the reader with a ranking of the sources based on our expectation that answers may be found in them.*

To focus on the core IR task, we assume that query and sources have already been translated to a temporally-tagged logical representation, e.g., using techniques from (Nguyen et al. 2015; LeBlanc and Balduccini 2016). We also assume the availability of suitable knowledge repositories (Suchanek et al. 2008; Inclezan 2016). It should be noted that, while our work is somewhat related to research on temporal relations (e.g., Allen's interval calculus), the two differ because we focus on reasoning about events and their effects, rather than relations between events.

The main contributions of this paper are (a) the exploration of a non-trivial variant of IR in which sources include sequences of events, and queries are about the state of the world after such events; (b) the extension of techniques for representing dynamic domains to increase the flexibility of the reasoning processes in the presence of uncertainty; (c) a formalization of the IR task based on action languages; (d) an automated IR procedure based on Answer Set Programming (ASP).

The paper begins with needed preliminaries. Next, we present a series of toy scenarios that guide the analysis of problem and reasoning processes. We formalize the reasoning task, present an ASP-based procedure for carrying it out automatically, and demonstrate it on selected scenarios. Finally, we briefly present related work, draw conclusions and discuss future work.

2 Preliminaries

In this paper, we build upon action language $\mathcal{AL}$ (Baral and Gelfond 2000) for the representation of knowledge about actions and their effects. The syntax of $\mathcal{AL}$ builds upon an alphabet consisting of a set $\mathcal{F}$ of symbols for *fluents* and a set $\mathcal{A}$ of symbols for *actions*. Fluents are boolean properties of the domain, whose truth value may change over time. A *fluent literal* is a fluent f or its negation $\neg f$. The statements of $\mathcal{AL}$ are:

$$a \text{ causes } l_0 \text{ if } l_1, l_2, \ldots, l_n \tag{1}$$

$$l_0 \text{ if } l_1, \ldots, l_n \tag{2}$$

$$a \text{ impossible\_if } l_1, \ldots, l_n \tag{3}$$

(1) is a *dynamic (causal) law*, and intuitively says that, if action a is executed in a state in which literals $l_1, \ldots, l_n$ hold, then l_0, the *consequence of the law*, will hold in the next state. (2) is a *state constraint* and says that, in any state in which $l_1, \ldots, l_n$ hold, l_0 also holds. (3) is an *executability condition* and says that a cannot be executed if $l_1, \ldots, l_n$ hold. A set of statements of $\mathcal{AL}$ is called *action description*. The semantics of $\mathcal{AL}$ maps action descriptions to transition diagrams. A set S of literals is *closed under a state constraint* (2) if $\{l_1, \ldots, l_n\} \not\subseteq S$ or $l_0 \in S$. S is *consistent* if, for every $f \in \mathcal{F}$, at most one of f, $\neg f$ is in S. It is *complete* if at least one of f, $\neg f$ is in S. A *state* of an action description AD is a complete and consistent set of literals closed under the state constraints of AD.

Given an action a and a state σ, the set of *(direct) effects of a in σ*, denoted by $E(a, \sigma)$, is the set that contains a literal l_0 for every dynamic law (1) such that $\{l_1, \ldots, l_n\} \subseteq \sigma$. Given a set S of extended literals and a set Z of state constraints, the *set*, $Cn_Z(S)$, *of consequences of S under Z* is the smallest set of extended literals that contains S and is closed under Z. Finally, an action a is *non-executable* in a state σ if there exists an executability condition (3) such that $\{l_1, \ldots, l_n\} \subseteq \sigma$. Otherwise, the action is *executable* in σ.

The semantics of an action description AD is defined by its *transition diagram* $\tau(AD)$, a directed graph $\langle N, E \rangle$ such that: N is the collection of all states of AD, and E is the set of all triples $\langle \sigma, a, \sigma' \rangle$ where σ, σ' are states, a is an action executable in σ, and σ' satisfies the *successor state equation* $\sigma' = Cn_Z(E(a, \sigma) \cup (\sigma \cap \sigma'))$, where Z is the set of all state constraints of AD. Triple $\langle \sigma, a, \sigma' \rangle$ is called a transition of $\tau(AD)$ and σ' is a *successor state of* σ (under a). A *path* in a transition diagram $\mathcal{T}(A)$ is a sequence $\langle \sigma_0, a_0, \sigma_1, a_1, \sigma_2, \ldots, \sigma_n \rangle$ in which every triple $\langle \sigma_i, a_i, \sigma_{i+1} \rangle$ satisfies the successor state equation. We denote the initial state of a path π by π_{σ_0}.

Next, we introduce ASP (Gelfond and Lifschitz 1991). Let Σ be a signature containing constant, function and predicate symbols. Terms and atoms are formed as in first-order logic. A *literal* is an atom a or its negation $\neg a$. A *rule* is a statement of the form: $h_1, \ldots, h_k \leftarrow l_1, \ldots, l_m,$ $\text{not } \; l_{m+1}, \ldots, \text{not } \; l_n$ where h_i's and l_i's are literals and *not* is called *default negation* operator. Its intuitive meaning in terms of a rational agent reasoning about its beliefs is "if you believe $\{l_1, \ldots, l_m\}$ and have no reason to believe $\{l_{m+1}, \ldots, l_n\}$, then you must believe one of $\{h_1, \ldots, h_k\}$." If $m = n = 0$, symbol $\leftarrow$ is omitted and the rule is a *fact*. Rules of the form $\bot \leftarrow l_1, \ldots, \text{not } l_n$ are abbreviated $\leftarrow l_1, \ldots, \text{not } l_n$, and called *constraints*, intuitively meaning that $\{l_1, \ldots, \text{not } l_n\}$ must not be satisfied. A rule with variables is interpreted as a shorthand for the set of rules obtained by replacing the variables with all possible variable-free terms. A *program* is a set of rules over Σ. A consistent set S of literals is closed under a rule if $\{h_1, \ldots, h_k\} \cap S \neq \emptyset$ whenever $\{l_1, \ldots, l_m\} \subseteq S$ and $\{l_{m+1}, \ldots, l_n\} \cap S = \emptyset$. Set S is an answer set of a *not*-free program Π if S is the minimal set closed under its rules. The reduct, Π^S, of a program Π w.r.t. S is obtained from Π by removing every rule containing an expression "not l" s.t. $l \in S$ and by removing every other occurrence of not l. Set S is an answer set of Π if it is the answer set of Π^S.

3 Problem Analysis

The previous example allows us to provide a first high-level characterization of the task we aim to study, as one in which we are given a query $\mathcal{Q}$ and a collection of sources $\mathcal{S}_1, \ldots, \mathcal{S}_n$, and are asked to produce scores $s_1, \ldots, s_n$ indicating how relevant each source is to the task of finding an answer to $\mathcal{Q}$. If we adopt the convention that 0 is the best possible score and ∞ the worst, then it is conceivable that, in Example 1, $\mathcal{S}_1$ should be assigned a score of 0 and $\mathcal{S}_2$ a score of ∞ to indicate complete irrelevance.

As in traditional Information Retrieval (IR), the sources will be ranked based on their respective score. We expect that, in the long-term, both syntactic and semantic aspects will have to be taken into considerations in determining scores for the documents. Thus, below, we use the term "semantic score" when referring to the score assigned to documents by the techniques we are studying. It is worth stressing the difference between the task at hand and question answering, where the goal is to produce a definitive answer. At the end of the process we consider here, the answer to $\mathcal{Q}$ may still be unknown, but there will be reason to believe that careful study by a human of the sources identified as relevant will lead to such answer.

Next, we consider a number of examples and corresponding expectations. Based on the examples, later we propose a formalization of the reasoning processes. Example 1 showed that the event of going on a first date may lead us to infer that John is not married. But how can one reach such conclusion? One option is to reason by cases, and consider two possible views of the world: one in which John is married at the beginning of the story, and one in which he is not. Commonsense tells us that the action[1] of going on a first date is not executable when married. Hence, the view in which John is initially married is inconsistent with the source. So, we conclude that John must not have been married in the initial state. Given further knowledge that one does not get married on a first date, one can infer that John remains not married after the date. Thus, the source provides evidence that a reader can use to answer the query.

From a technical perspective, the example highlights the importance of being able to reason by cases, to reason about the executability of actions, and to propagate the truth of properties of interest over the duration of the story. Note, however, that reasoning by cases is sometimes misleading. Consider $\mathcal{S}_2$ from Example 1: reasoning by cases leads to the same two possible initial states. Since reading does not affect married status, there are two ending states for the story. This might be taken as an indication that the source provides *some* useful evidence for a

[1] From now on, we will use action and event as synonyms.

reader, but it is clear intuitively that $\mathcal{S}_2$ is, in fact, irrelevant. Next, let us consider if, and how, the previous query should match a more complex document. For the sake of this example, let us assume the existence of a fictitious country C, whose laws allow plural marriage.

Example 2
$\mathcal{Q}$: Is John married?
$\mathcal{S}$: John, who lives in country C, just went on his first date with Mary.

In this case, $\mathcal{S}$ does not provide useful information towards answering $\mathcal{Q}$. John is from C, where plural marriage is allowed, and knowledge about plural marriage yields that being married does not preclude a married person from going on a first date. The example also demonstrates the importance of reasoning about *default statements* (statements that are *normally* true) and their exceptions. The fact that, normally, married people do not go on first dates is an instance of a default statement, and an inhabitant of C constitutes an exception to it. Similarly to $\mathcal{S}_2$ from the previous example, reasoning by cases may be somewhat misleading, as it may suggest that the source provides some evidence useful to answering the question. Rather than reasoning by cases, it appears to be more appropriate to state that whether John is initially married is *unknown*. The lack of knowledge is propagated to the final state, given that going on a date has no effect on it in the present context. The source is thus irrelevant and should receive a semantic score of ∞. Note the striking difference in scores between $\mathcal{S}_1$ from the previous example and the current source: it appears that in some cases reasoning by cases is useful, while in others reasoning explicitly about lack of knowledge is more appropriate. In the next section, we provide a characterization of reasoning matching this intuition. Next, we investigate the role of the effects of actions.

Example 3
$\mathcal{Q}$: Is John married?
$\mathcal{S}$: John, who lives in country C, recently went on his first date with Mary. A week later, they tied the knot in Las Vegas.

Obviously, a first indication of relevance can be obtained with shallow reasoning and syntactic matching: "tying the knot" is a synonym of "getting married" and "getting married" and "being married" share enough similarities to make a match likely. However, we are interested in more sophisticated reasoning. In the initial state, John may or may not be married due to his country's laws. Similarly to Example 1, John's married status persists in the state following the first date. Tying the knot, however, has the effect of making John married in the resulting state. Hence, $\mathcal{S}$ is indeed relevant to $\mathcal{Q}$. Intuitively, its semantic score should be equal to that of $\mathcal{S}_1$ from Example 1. This demonstrates the importance of keeping track of the changes in the truth of the relevant properties over time. The next example takes this argument one step further.

Example 4
$\mathcal{Q}$: Is John married?
$\mathcal{S}$: John recently went on his first date with Mary. A week later, they tied the knot in Las Vegas. A month later, they filed for divorce.

Here, we assume that filing for divorce does not immediately cause the spouses to be divorced. For simplicity of presentation, we adopt a view in which filing for divorce has a non-deterministic effect: in the resulting state, it is equally likely for the spouses to be married or not. The relevance of $\mathcal{S}$ to $\mathcal{Q}$ is not as straightforward as in some of the previous cases. It is indeed true that, at the end of the story, it is unknown whether John is married. On the other hand, the story still provides some information pertaining to John's married status – certainly, more than source $\mathcal{S}_2$ ("John read a book") from Example 1 or the source from Example 2 ("John, who lives in country C, just went on his first date with Mary.").

One way to make a distinction between the two cases is to consider that, if $\mathcal{S}$ from Example 4 is provided to a reader, and the reader manages to determine whether the filing action succeeded (e.g., by gathering additional evidence), $\mathcal{S}$ will immediately allow the reader to answer $\mathcal{Q}$. Differently from the previous examples, knowing that filing occurred is *essential* to allowing a reader to answer the question. In conclusion, while $\mathcal{S}$ is not as relevant to $\mathcal{Q}$ as other sources we have considered, it is still somewhat relevant. This will have to be reflected in the score assigned to the source, which should be higher than the 0 assigned to $\mathcal{S}_1$, but obviously smaller than ∞ because the source is indeed relevant. Next, we propose a formalization that captures the behaviors described.

4 Formalization of the Reasoning Task

Our formalization leverages techniques from the research on reasoning about actions and change, and specifically action language $\mathcal{AL}$ (Baral and Gelfond 2000), approximated representations (Morales et al. 2007) and evidence-based reasoning (Balduccini and Gelfond 2003). These techniques rely on a graph-based representation of the evolution of the state of the world over time in response to the occurrence of actions. We adopt and expand this approach. Specifically, similarly to (Morales et al. 2007), our formalization enables reasoning explicitly about lack of knowledge. Differently from it, however, we allow a reasoner to reason by cases whenever needed. This is applied to knowledge about both initial state and effects of actions. Our approach also leverages evidence-based reasoning to rule out some of the cases considered. Finally, we adopt $\mathcal{AL}$ as the underlying formalism, but expand it for an explicit characterization of non-deterministic effects and we allow hypothesizing about exceptional/atypical circumstances, eventually linking them to the relevance of sources. Differently from $\mathcal{AL}$, our language is defined so that, in the presence of actions with non-deterministic effects, it is possible to reason both by cases, and by explicitly characterizing lack of knowledge. The syntax of the resulting language, which we call $\mathcal{AL}_{IR}$, is described next by building on that of $\mathcal{AL}$, followed by its semantics.

In $\mathcal{AL}_{IR}$, we identify a (possibly empty) subset $\mathcal{D}$ of $\mathcal{F}$ called the set of *default fluents*. Default fluents are assumed false at the beginning of a sequence of events. Additionally, an *extended (fluent) literal* is either a fluent literal or the expression $u(f)$, intuitively meaning that it is unknown whether f is true or false. Expression $u(f)$ is called *proper extended literal*. The syntax of dynamic law (1) is extended to allow l_0 to be an proper extended literal. If l_0 is a proper extended literal $u(f)$, the law intuitively states that the action affects the truth of f non-deterministically. The action of filing for divorce from Example 4 might be modeled with a dynamic law that has $u(married)$ as its consequence.

The semantics of $\mathcal{AL}_{IR}$ is obtained by extending the definitions to extended literals as needed. Specifically, a set S of extended literals is consistent if, for every $f \in \mathcal{F}$, at most one of f, $\neg f$, $u(f)$ is in S. It is *complete* if at least one of f, $\neg f$, $u(f)$ is in S. A *state* of an action description AD of $\mathcal{AL}_{IR}$ is a complete and consistent set of extended literals closed under the state constraints of AD.

In this phase of the investigation, we restrict our attention to cases in which every action has at most a single direct non-deterministic effect, and we disallow concurrent actions. Lifting these restrictions is not difficult, but complicates the presentation. The direct effects of actions are extended as follows. Given an action a and a state σ, the set of *combined (direct) effects of* a *in* σ, denoted by $E(a^?, \sigma)$, coincides with $E(a, \sigma)$ from $\mathcal{AL}$. The set of *positive effects of* a *in* σ, $E(a^+, \sigma)$, is the set that contains: (a) a fluent literal l for every dynamic law (1) such that $l_0 = l$ and $\{l_1, \dots, l_n\} \subseteq \sigma$, and (b) a fluent f for every dynamic law such that $l_0 = u(f)$ and $\{l_1, \dots, l_n\} \subseteq \sigma$. Similarly, the set of *negative effects of* a *in* σ, $E(a^-, \sigma)$, is the set that contains: (a) a fluent literal l for every dynamic law such that $l_0 = l$ and $\{l_1, \dots, l_n\} \subseteq \sigma$, and (b) a fluent literal $\neg f$ for every dynamic law such that $l_0 = u(f)$ and $\{l_1, \dots, l_n\} \subseteq \sigma$.

Given an action description AD, the edges of the corresponding transition diagram are given by all triples $\langle \sigma, a^\circ, \sigma' \rangle$ where σ, σ' are states, a is an action executable in σ, $\circ$ is one of $?, +, -$, and σ' satisfies the equation:

$$\sigma' = Cn_Z(E(a^\circ, \sigma) \cup (\sigma \cap \sigma')).$$

When multiple successor states exist for a given σ and a°, the action description is called *non-deterministic*.

A dynamic law with a proper extended literal $u(f)$ as its consequence has two *deterministic counterparts*, obtained by replacing its consequence by f and $\neg f$ respectively. A dynamic law with a fluent literal as its consequence has a single deterministic counterpart, which coincides with the law itself. An action description AD has *emergent non-deterministic behavior* if there exists a non-deterministic action description AD', obtained from AD by replacing every dynamic law by one of its deterministic counterparts. In the current phase of the investigation, we do not consider action descriptions with emergent non-deterministic behavior.[2]

Next, we turn our attention to the use of transition diagrams to reason about sequences of actions and to determine the relevance of available sources.

5 Reasoning about Relevance of Sources

In our approach, a *qualified action sequence* is a tuple $s = \langle a_0/q_0, a_1/q_1, \dots, a_k/q_k \rangle$ where a_i's are actions and each q_i is one of $?$, $\times$. Intuitively, qualifier $?$ specifies that the combined effects of the action should be considered, while $\times$ indicates that reasoning by cases should be used. The *length of* s is $k+1$. The *degree of* s, denoted by $|s|$, is the number of expressions of the form $a_i/\times$ in s. If $\aleph = \langle a_0, a_1, \dots, a_k \rangle$ is a sequence of actions, we say that $s = \langle a_0/q_0, a_1/q_1, \dots a_k/q_k \rangle$ *extends* $\aleph$ for every possible choice of qualifiers. $\aleph^?$ denotes the extension of $\aleph$ where all qualifiers are $?$ and $\aleph^\times$ denotes the extension where all qualifiers are $\times$. Let σ be a state and s be a qualified action sequence. A path $\pi = \langle \sigma_0, \alpha_0, \sigma_1, \dots, \alpha_k, \sigma_{k+1} \rangle$ is a *model of* σ, s if all of the following hold: (a) $\sigma_0 = \sigma$, (b) if $q_i =?$, then $\alpha_i = a_i^?$, (c) if $q_i = \times$, then $\alpha_i = a_i^+$ or $\alpha_i = a_i^-$.

Given a set Σ of states and a qualified action sequence s, a path π is a *model of* Σ, s if π is a model of σ, s for some $\sigma \in \Sigma$. To illustrate these notions, consider an action description $\{a_1$ causes $\neg g$ if g; a_2 causes $u(f)$ if $\neg g\}$. Let σ be $\{\neg f, g\}$. It is not difficult to see that the pair $\sigma, \langle a_1/?, a_2/? \rangle$ has a unique model, $\langle \{\neg f, g\},\ a_1^?,\ \{\neg f, \neg g\},\ a_2^?,\ \{u(f), \neg g\} \rangle$. On the other hand, $\sigma, \langle a_1/?, a_2/\times \rangle$ has two models, $\langle \{\neg f, g\},\ a_1^?,\ \{\neg f, \neg g\},\ a_2^+,\ \{f, \neg g\} \rangle$ and $\langle \{\neg f, g\}, a_1^?, \{\neg f, \neg g\},\ a_2^-, \{\neg f, \neg g\} \rangle$. The degrees of the two qualified action sequences are 0 and 1 respectively.

Let us now consider cases in which knowledge about the initial state is incomplete. Intuitively, if the truth value of f is unknown, one may assume that f is false if it is a default fluent and that

[2] Action description $\{q$ if $\neg r, p;\ r$ if $\neg q, p; a$ causes $p\}$ has an emergent non-deterministic behavior.

$u(f)$ holds otherwise. However, as highlighted in the above examples, it is sometimes necessary to consider other options for certain fluents. This intuition is captured by the notion of forcing of a fluent. Given a consistent set I of extended literals and a fluent f, $I[f]$ denotes the set $\mathcal{I}$ defined as follows, called the *forcing of* f *in* I: if $f \in \mathcal{D}$ and $\{\neg f, u(f)\} \cap I = \emptyset$, then $\mathcal{I} = \{I \cup \{f\}\}$; if $f \notin \mathcal{D}$ and $\{f, \neg f, u(f)\} \cap I = \emptyset$, then $\mathcal{I} = \{\ I \cup \{f\},\ I \cup \{\neg f\}\ \}$; otherwise, $\mathcal{I} = \{I\}$. For sets of fluents, the *forcing of* $\{f_1, \ldots, f_m\}$ *in* I, written $I[\{f_1, \ldots, f_m\}]$, is defined as follows: (a) if $m = 1$, then $I[\{f_1\}] = I[f_1]$; (b) if $m > 1$, then $I[\{f_1, \ldots, f_m\}] = \{I'[f_m] \mid I' \in I[\{f_1, \ldots, f_{m-1}\}]\}$.

As an example, let us apply these definitions to $\mathcal{S}_1$ from Example 1, "John went on his first date with Mary." Assume that the translation from natural language yields[3] $\mathcal{Q} = m$, $\mathcal{F} = \{m, ab\}$, $\mathcal{D} = \{ab\}$, $I = \emptyset$ and $\aleph = \langle d \rangle$. Let us also assume that the action description, AD, is $\{$impossible d if $m, \neg ab\}$.[4] Note the use of default fluent ab to formalize the fact that the action is *normally* impossible if one is married. It is not difficult to see that $I[\mathcal{F} \setminus \mathcal{D}] = I[\{m, ab\} \setminus \{ab\}]$ is $\{\{m\}, \{\neg m\}\}$, indicating that, in the initial state, we can assume that he may or may not have been married.

Let Z be the set of state constraints of AD. The *default closure* of I is the set $\delta(I) = Cn_Z(I \cup \{\neg f \mid f \in \mathcal{D} \wedge f \notin I\})$. If $\delta(I)$ is consistent, we say that the *completion of* I is the set of extended literals $\gamma(I) = \delta(I) \cup \{u(f) \mid f \notin \delta(I) \wedge \neg f \notin \delta(I)\}$. Note that $\gamma(I)$ may not exist, as in the case of $I = \{p, q\}$ and of $AD = \{\neg q$ if $p\}$. If $\gamma(I)$ exists, it is complete, consistent and includes I. Given a set F of fluents, the *completion of* I *w.r.t.* F is the set $\gamma(I, F) = \{\gamma(I') \mid I' \in I[F] \wedge \gamma(I')$ exists$\}$. The *degree of* $\gamma(I, F)$, denoted by $|\gamma(I, F)|$, is $|F|$.

Going back to Example 1, applying the closure to each element of $I[\mathcal{F} \setminus \mathcal{D}]$ yields, respectively, $\{m, \neg ab\}$ and $\{\neg m, \neg ab\}$, which can intuitively be viewed as the initial states that are consistent with assumptions made about m.

As demonstrated by Example 1, there are cases in which the truth of certain fluents in the initial state can be inferred indirectly from the source. The following definition of $\rho(I, \aleph)$ captures this idea. Given a consistent set I of extended fluent literals and a sequence of actions $\aleph$:

$$\rho(I, \aleph) = \bigcap_{I' \in I[\mathcal{F} \setminus \mathcal{D}]} \{I' \mid \gamma(I'), \aleph^{\times} \text{ has a model}\}$$

Note that $\rho(I, \aleph)$ may not exist, e.g., if $\gamma(I')$ does not exist for any element of $I[\mathcal{F} \setminus \mathcal{D}]$. If $\rho(I, \aleph)$ does not exist, then the source is irrelevant and its semantic score if ∞. If, instead, $\rho(I, \aleph)$ exists, it is not difficult to see that $I \subseteq \rho(I, \aleph)$.

Let us see how $\rho(I, \aleph)$ is calculated in Example 1. The first step consists in checking for models of $\gamma(I'), \aleph^{\times}$. Clearly, $\{m, \neg ab\}, \langle d \rangle$ has no model, because d is not executable. On the other hand, $\{\neg m, \neg ab\}, \langle d \rangle$ has a model. Hence, $\rho(I, \aleph)$ is the intersection of the only set $\{\neg m\}$, resulting in $\rho(I, \aleph) = \{\neg m\}$. Intuitively, this mirrors the intuition that John is not married in the initial state.

We are now ready to introduce the notion of entailment and to use it to determine whether there is a match between $\mathcal{Q}$ and $\mathcal{S}$. A path $\pi = \langle \sigma_0, \alpha_0, \sigma_1, \ldots, \alpha_{k-1}, \sigma_k \rangle$ *entails* a fluent literal l (written $\pi \models l$) if $l \in \sigma_k$. Given a fluent f, we say that π *entails* $\pm f$ (written $\pi \models \pm f$) if $\pi \models f$ or $\pi \models \neg f$.

[3] We use abbreviations to save space. Fluents: m – John is married; ab – John is an exception w.r.t. going on first dates when married. Actions: d – going on a first date; r – reading a book.

[4] In practice, variables may be introduced to increase generality.

For simplicity, we assume $\mathcal{Q}$ to be a fluent. Let I be a set of fluent literals explicitly stated to hold in the initial state by $\mathcal{S}$ and let $\aleph = \langle a_0, a_1, \ldots, a_k \rangle$ be the sequence of actions from $\mathcal{S}$. We say that $\mathcal{S}$ *is a match for* $\mathcal{Q}$ if there exist a set F of fluents and a qualified action sequence s extending $\aleph$ s.t.:

c1 π entails $\pm\mathcal{Q}$ for some model π of $\gamma(\rho(I, \aleph), F), s$, and
c2 for every model π' of $\gamma(\pi_{\sigma_0} \setminus \rho(I, \aleph), \emptyset), \langle\, \rangle$, one of the following holds: (a) $\pi' \not\models \pm\mathcal{Q}$, or (b) $\pi' \models \neg\mathcal{Q}$ and $\pi \models \mathcal{Q}$, or (c) $\pi' \models \mathcal{Q}$ and $\pi \models \neg\mathcal{Q}$.

Intuitively, the first condition checks whether the document is relevant to the query – possibly under some assumptions about the default fluents – while the second condition ensures that such assumptions are not directly and solely responsible for the fact that the document is relevant.

The semantic score of $\mathcal{S}$ is the smallest value of $|\gamma(\rho(I, \aleph), F)| + |s|$ for all possible choices of F and s satisfying the above items. If no F and s were found to satisfy the above conditions, then $\mathcal{S}$ is not a match for $\mathcal{Q}$ (i.e., it is irrelevant to the query) and its semantic score is ∞.

In reference to Example 1, let us first look for F, s, satisfying (**c1**). Let us begin with $F = \emptyset$, $s = \langle d \rangle^?$, which have a degree of 0. It is not difficult to see that $\gamma(\rho(I, \aleph), F) = \gamma(\{\neg m\}, \emptyset)\rangle = \{\{\neg m, \neg ab\}\}$ and that $\{\{\neg m, \neg ab\}\}, \langle d \rangle^?$ has a unique model $\pi = \langle \{\neg m, \neg ab\}, d^?, \{\neg m, \neg ab\} \rangle$. Thus, the model entails $\pm Q$, which means that condition (**c1**) for establishing a match is satisfied.

Next, we check condition (**c2**). Clearly, $\gamma(\pi_{\sigma_0} \setminus \rho(I, \aleph), \emptyset) = \{\{u(m), \neg ab\}\}$. $\{\{u(m), \neg ab\}\}$, $\langle\rangle$ has a unique model, $\langle \{u(m), \neg ab\} \rangle$, and it does not entail $\pm Q$. Intuitively, this means that the assumption made about the initial state *is not directly responsible* for the ability to entail the query in (**c1**). Hence, $\mathcal{S}$ matches $\mathcal{Q}$. Additionally, because $F = \emptyset$, $s = \langle d \rangle^?$ yield a score of 0, the semantic score of the document is 0.

As an additional example, consider $\mathcal{S}_2$, "John read a book," from Example 1. As above, $\mathcal{Q} = m$, $\mathcal{F} = \{m, ab\}$, $\mathcal{D} = \{ab\}$, and $I = \emptyset$, while and $\aleph = \langle r \rangle$. AD is the same as before.[5] $I[\mathcal{F} \setminus \mathcal{D}]$ is $\{\{m\}, \{\neg m\}\}$, yielding a closure of $\{\{m, \neg ab\}, \{\neg m, \neg ab\}\}$. This time, both $\{m, \neg ab\}, \langle\, r \rangle^\times$ and $\{\neg m, \neg ab\}, \langle\, r \rangle^\times$ have models. Hence, $\rho(I, \aleph) = \{m\} \cap \{\neg m\} = \emptyset$. That is, the initial truth value of no fluent can be inferred from the story.

Next, we consider the models of $\gamma(\rho(I, \aleph), F), s$. Consider $F = \emptyset$, $s = \langle r \rangle^?$, with a degree of 0. $\gamma(\rho(I, \aleph), F) = \gamma(\emptyset, \emptyset)\rangle = \{\{u(m), \neg ab\}\}$. $\{\{u(m), \neg ab\}\}, \langle r \rangle^?$ has a unique model $\pi = \langle \{u(m), \neg ab\}, r^?, \{u(m), \neg ab\} \rangle$. Clearly, $\pi \not\models \pm Q$.

The next possible options, with a combined degree of 1, are $F = \emptyset$, $s = \langle r \rangle^\times$ and $F = \{m\}$, $s = \langle r \rangle^?$. In the first case, there are two models, e.g., $\pi = \langle \{u(m), \neg ab\}, r^+, \{u(m), \neg ab\} \rangle$, but neither entails $\pm Q$. The second case is more interesting. Clearly, there are two models of $\gamma(\rho(I, \aleph), F), s = \gamma(\emptyset, \{m\}), \langle r \rangle^?$: $\pi = \langle \{m, \neg ab\}, r^?, \{m, \neg ab\} \rangle$ and $\pi' = \langle \{\neg m, \neg ab\}, r^?, \{\neg m, \neg ab\} \rangle$, and $\pi \models Q$, while $\pi' \models \neg Q$. Hence, we need to check condition (**c2**) for each. For the former, $\gamma(\pi_{\sigma_0} \setminus \emptyset, \emptyset) = \{\{m, \neg ab\}\}$, and $\{\{m, \neg ab\}\}, \langle\rangle$ has a unique model $\langle \{m, \neg ab\} \rangle$, which entails Q. Thus, the condition is not satisfied. For π', we obtain a unique model $\langle \{\neg m, \neg ab\} \rangle$, which entails $\neg Q$, failing to satisfy the condition as well. Therefore, none of these choices for F and s yields a match. Similar conclusions can be drawn for the other choices for F and s. Hence, $\mathcal{S}_2$ does not match Q and receives a semantic score of ∞. The other examples are solved similarly. The details are omitted to save space, but we provide highlights of some of them.

Example 2. Contrast the previous case with Example 2. People from countries that allow plural marriage are exceptions to the custom about first dates, and thus $I = \{ab\}$, $\aleph = \langle d \rangle$, and

[5] We oversimplify the action description for sake of clarity.

$I[\mathcal{F} \setminus \mathcal{D}] = \{\{m, ab\}, \{\neg m, ab\}\}$. Differently from the previous case, both sets of $I[\mathcal{F} \setminus \mathcal{D}]$ yield a model, since ab makes the executability condition inapplicable. Hence, $\rho(I, \aleph) = \{ab\}$. Selecting $F = \emptyset$, $s = \langle d \rangle^?$ yields a unique model $\langle \{u(m), ab\}, d^?, \{u(m), ab\} \rangle \not\models \pm Q$. Selecting $F = \{m\}$, $s = \langle d \rangle^?$ yields two models entailing Q and $\neg Q$ respectively, but the same are entailed by $\gamma(\pi_{\sigma_0} \setminus \rho(I, \aleph), \emptyset), \langle\rangle$, thus failing condition (**c2**). Similar reasoning applies to the other cases. Because no F, s could be identified, the semantic score of $\mathcal{S}$ is ∞, indicating that it is irrelevant to $\mathcal{Q}$. Note the key role played by condition (**c2**) in this example: without it, the source would have been deemed relevant to the query.

<u>Example 4.</u> Consider Example 4, where the action description is expanded with $\{w$ causes m; fd causes $u(m)\}$ and relevant executability conditions. We have $I = \emptyset$, $\aleph = \langle d, w, fd \rangle$, and, similarly to Example 1, $\rho(I, \aleph) = \{\neg m\}$. The model obtained from $F = \emptyset$, $s = \aleph^?$ does not entail $\pm Q$. On the other hand, $F = \emptyset$, $s = \langle d/?, w/?, fd/\times \rangle$, yield two models, entailing Q and $\neg Q$ resp., depending on the outcome of fd. This time, condition (**c2**) is satisfied, since, in both cases, $\gamma(\pi_{\sigma_0} \setminus \rho(I, \aleph), \emptyset) = \{\{u(m), \neg ab\}\}$ and $\{\{u(m), \neg ab\}\}, \langle\rangle$ does not entail $\pm Q$. In conclusion, $\mathcal{S}$ indeed matches $\mathcal{Q}$, and the source has semantic score $|\emptyset| + |\langle d/?, w/?, fd/\times \rangle| = 1$. As expected, its semantic score is worse than that of, e.g., $\mathcal{S}_1$, while obviously better than that of, e.g., $\mathcal{S}_2$.

6 Automating the Reasoning Task

Next, we automate the reasoning task discussed earlier by means of a translation of $\mathcal{AL}_{IR}$ to ASP. Given a set I of extended fluent literals, a set F of fluents, a qualified action sequence s, and an action description AD, the encoding of $\mathcal{AL}_{IR}$ is program $\Pi_{AD}(I, F, s)$, described next.

In the following, $\mathtt{I}$ ranges over steps in the evolution of the domain[6]; given fluent literal l, $\chi(l, \mathtt{I})$ stands for $holds(f, \mathtt{I})$ if $l = f$ and $\neg holds(f, \mathtt{I})$ if $l = \neg f$. For every action a, the translation includes a rule $pos(a, \mathtt{I}) \vee neg(a, \mathtt{I}) \leftarrow occurs(a, \mathtt{I}), split(a, \mathtt{I})$. The translation of a dynamic law (1) depends on the form of l_0. If l_0 is a fluent literal, translation is: $\chi(l_0, \mathtt{I}+1) \leftarrow occurs(a, \mathtt{I}), \chi(l_1, \mathtt{I}), \ldots, \chi(l_n, \mathtt{I})$. If l_0 is of the form $u(f)$, the translation of the law is:

$$u(f, \mathtt{I}+1) \leftarrow occurs(a, \mathtt{I}), \chi(l_1, \mathtt{I}), \ldots, \chi(l_n, \mathtt{I}), \text{not } split(a, \mathtt{I}).$$
$$\chi(f, \mathtt{I}+1) \leftarrow pos(a, \mathtt{I}), \chi(l_1, \mathtt{I}), \ldots, \chi(l_n, \mathtt{I}).$$
$$\chi(\neg f, \mathtt{I}+1) \leftarrow neg(a, \mathtt{I}), \chi(l_1, \mathtt{I}), \ldots, \chi(l_n, \mathtt{I}).$$

Expression $occurs(a, \mathtt{I})$ states that action a occurs at step $\mathtt{I}$ in the story; $split(a, \mathtt{I})$ states that reasoning by cases should be applied to the outcomes of that occurrence of a. A state constraint (2) is translated as an ASP rule of the form $holds(l_0, \mathtt{I}) \leftarrow holds(l_1, \mathtt{I}), \ldots, holds(l_n, \mathtt{I})$. Executability condition (3) is translated as a rule $\leftarrow occurs(a, \mathtt{I}), \chi(l_1, \mathtt{I}), \ldots, \chi(l_n, \mathtt{I})$. The translation of an action description is completed by the inertia axioms, which are expanded in $\mathcal{AL}_{IR}$ to accommodate extended literals ($\mathtt{F}$ is a variable ranging over all fluents):

$$\chi(\mathtt{F}, \mathtt{I}+1) \leftarrow \chi(\mathtt{F}, \mathtt{I}), \text{not } \chi(\neg\mathtt{F}, \mathtt{I}+1), \text{not } u(\mathtt{F}, \mathtt{I}+1).$$
$$\chi(\neg\mathtt{F}, \mathtt{I}+1) \leftarrow \chi(\neg\mathtt{F}, \mathtt{I}), \text{not } \chi(\mathtt{F}, \mathtt{I}+1), \text{not } u(\mathtt{F}, \mathtt{I}+1).$$
$$u(\mathtt{F}, \mathtt{I}+1) \leftarrow u(\mathtt{F}, \mathtt{I}), \text{not } \chi(\mathtt{F}, \mathtt{I}+1), \text{not } \chi(\neg\mathtt{F}, \mathtt{I}+1).$$

The next axioms define the completion of the initial state:

$$[\mathbf{g_1}] \;\; \chi(\mathtt{F}, 0) \leftarrow init(\mathtt{F}). \qquad \chi(\neg\mathtt{F}, 0) \leftarrow \neg init(\mathtt{F}).$$

[6] We assume that the range of $\mathtt{I}$ is provided by the process of translating the passage to a logical representation.

$$
\begin{aligned}
[\mathbf{g_2}]\ \ & \chi(\mathtt{F},0) \leftarrow forced(\mathtt{F}), default(\mathtt{F}), \text{not } \neg init(\mathtt{F}).\\
& \chi(\mathtt{F},0) \vee \chi(\neg\mathtt{F},0) \leftarrow forced(\mathtt{F}), \text{not } default(\mathtt{F}),\\
& \qquad\qquad\qquad\qquad\qquad \text{not } init(\mathtt{F}), \text{not } \neg init(\mathtt{F}).\\
[\mathbf{g_3}]\ \ & \chi(\neg\mathtt{F},0) \leftarrow default(\mathtt{F}), \text{not } \chi(\mathtt{F},0).\\
& u(\mathtt{F},0) \leftarrow \text{not } default(\mathtt{F}), \text{not } \chi(\mathtt{F},0), \text{not } \chi(\neg\mathtt{F},0).
\end{aligned}
$$

Above, statement $default(f)$, included as fact for every $f \in \mathcal{D}$, states that f is a default fluent. $init(f)$ (resp., $\neg init(f)$) says that f is initially true (resp., false). $forced(f)$ states that f is part of a forcing. Rules $[\mathbf{g_1}]$ map the knowledge about the initial state to statements $holds(\cdot,\cdot)$. $[\mathbf{g_2}]$ formalizes to the notion of forcing. $[\mathbf{g_3}]$ defines the completion.

The next step of the definition of $\Pi_{AD}(I,F,s)$ is the encoding of its arguments. For every $f \in I$ (resp., $\neg f \in I$), $\Pi_{AD}(I,F,s)$ includes a fact $init(f)$ (resp., $\neg init(f)$). For every $f \in F$, $\Pi_{AD}(I,F,s)$ includes a fact $forced(f)$. Qualified action sequence s is encoded by a set of facts of the form $occurs(a,i)$ and $split(a,i)$, where a are actions from s and i are their indexes. Specifically, $a^?$ is translated as a statement $occurs(a,i)$, where i is the index in the sequence, while $a^\times$ is translated as two facts, $occurs(a,i)$, $split(a,i)$.

This completes the definition of $\Pi_{AD}(I,F,s)$. Next, we link its answer sets to the models of $\gamma(I,F), s$. We say that an answer set A *encodes a path* π if: (a) for every fluent literal l, $l \in \sigma_i$ iff $\chi(l,i) \in A$; (b) for every fluent f, $u(f) \in \sigma_i$ iff $u(f,i) \in A$; (c) for every action a, $\alpha_i = a^?$ iff $occurs(a,i) \in A$ and $split(a,i) \notin A$; (d) for every action a, $\alpha_i = a^+$ iff $\{occurs(a,i), split(a,i), pos(a,i)\} \subseteq A$; (e) for every action a, $\alpha_i = a^-$ iff $\{occurs(a,i), split(a,i), neg(a,i)\} \subseteq A$. The link is established by:

Proposition 1
Let I be a consistent set of fluent literals, F be a set of fluents, and s be a qualified action sequence. A path π is a model of $\gamma(I,F)$, s iff there exists an answer set of $\Pi_{AD}(I,F,s)$ that encodes π.

Corollary 1
A model π of $\gamma(I,F)$, s that entails l exists iff there exists an answer set A of $\Pi_{AD}(I,F,s)$ such that $\chi(l,k) \in A$, where k is the length of s. Also, for every fluent f, $\pi \models \pm f$ iff $\{\chi(f,k), \chi(\neg f,k)\} \cap A \neq \emptyset$.

These results motivate the algorithm in Figure 1. Let $||A||$ be the number of atoms of A formed by relations $forced$ and $split$. The behavior of the algorithm is characterized by:

Theorem 1
If $\mathcal{S}$ is a fluent, then $\mathcal{S}$ is a match for $\mathcal{Q}$ iff `FindMatch`(I,$\aleph$,$\mathcal{Q}$)$\neq \perp$. The rank of $\mathcal{S}$ is $||$`FindMatch`(I,$\aleph$,$\mathcal{Q}$)$||$.
Proof (sketch). Using the two previous results, the thesis is easily obtained by observing that step 1 implements the calculation of $\rho(I,\aleph)$, and that steps 4 and 4b check, respectively, conditions (**c1**) and (**c2**).

Let us trace the key parts of the algorithm with $\mathcal{S}_1$ from Example 1. Clearly, $\Pi_{AD}(I, \mathcal{F} \setminus \mathcal{D}, \aleph^\times\} \supseteq \{\leftarrow occurs(d,\mathtt{I}), holds(m,\mathtt{I}), step(\mathtt{I}).\ forced(m).\ occurs(d,0).\}$. Step 1 infers the initial truth of fluents indirectly from the $\mathcal{S}_1$, resulting in an answer set containing $\{\neg holds(m, 0),\ forced(m)\}$, i.e., John cannot be initially married. Hence, $I' = I \cup \{\neg m\}$. Step 4 checks condition (**c1**). It results in a unique answer set $A \supseteq \{holds(m,0), \neg holds(ab,0), occurs(d,0), \neg holds(m,1), \neg holds(ab,1)\}$, indicating that $\langle\{\neg m, \neg ab\}, d^?, \{\neg m, \neg ab\}\rangle$ entails $\pm m$. Step

Algorithm: `FindMatch`(I,$\aleph$,$\mathcal{Q}$)
Input: I – (set) fluent literals explicitly stated to hold in the initial state by $\mathcal{S}$; $\aleph = \langle a_0, a_1, \dots, a_k \rangle$ – sequence of actions from $\mathcal{S}$; $\mathcal{Q}$ – fluent.
Output: an answer set encoding a path if a match exists; $\bot$ otherwise.

1. Let R be the intersection of all answer sets of $\Pi_{AD}(I, \mathcal{F} \setminus \mathcal{D}, \aleph^{\times})$ and I' be $I \cup \{l \mid \{\chi(l,0), forced(f)\} \subseteq R \wedge (l = f \vee l = \neg f)\}$.
2. If $\Pi_{AD}(I, \mathcal{F} \setminus \mathcal{D}, \aleph^{\times})$ has no answer set, return $\bot$ and terminate.
3. Initialize $F := \emptyset$ and $s := \aleph^{?}$.
4. For every answer set A of $\Pi_{AD}(I', F, s)$ such that $\{\chi(\mathcal{Q}, k+1), \chi(\neg\mathcal{Q}, k+1)\} \cap A \neq \emptyset$:
   (a) Let $X = \{f \mid holds(f,0) \in A \wedge f \notin I'\} \cup \{\neg f \mid \neg holds(f,0) \in A \wedge \neg f \notin I'\}$.
   (b) For every answer set B of $\Pi_{AD}(X, \emptyset, \langle\,\rangle)$, check that $\{\chi(\mathcal{Q},0), \chi(\neg\mathcal{Q},0)\} \cap B = \emptyset$, or $\chi(\mathcal{Q},0) \in B \wedge \chi(\neg\mathcal{Q}, k+1) \in A$, or $\chi(\neg\mathcal{Q},0) \in B \wedge \chi(\mathcal{Q}, k+1) \in A$.
   (c) If every B satisfies the condition, then return A and terminate.
5. Select a set F' of fluents and an extension s' of $\aleph$ such that:
   (a) the pair F', s' has not yet been considered by the algorithm, and
   (b) $|F'| + |s'|$ is minimal among such pairs.
6. If no such pair F', s' exists, then return $\bot$ and terminate.
7. $F := F'$; $s := s'$. Repeat from step 4.

Fig. 1. `FindMatch` algorithm

4b checks condition (**c2**). There is a single answer set $B \supseteq \{u(m,0), \neg holds(ab,0), u(m,1), \neg holds(ab,1)\}$, and, clearly, $\{holds(m,0), \neg holds(m,0)\} \cap B = \emptyset$. Hence, (**c2**) is satisfied and the algorithm returns A. The rank of $\mathcal{S}_1$ is $||A|| = 0$.

7 Related Work

The IR task (Korfhage 1997) aims at identifying, among a set of available documents, those that are most relevant to a query provided by the user. In the traditional IR approach to representing documents, the text is fragmented into lists of keywords, terms, and other content descriptors. When presented with a query, an IR system determines the relevance of a document to the query by measuring the overlap of terms between the query and a particular document (Manning et al. 2008). Most IR systems base the relevance of a document on a syntactic measurement of the overlap of terms between query and document (Manning et al. 2008). Results using this approach are improved via the application of query expansion (Carpineto and Ramano 2012), an approach that reformulates the original query to expand the sphere of search, for example by collecting synonyms for terms in the query and searching for documents related to those synonyms.

A number of approaches have been proposed to improve search results. A recent approach (Blanco and Lioma 2012) aims to rethink the modeling of documents by representing text as a graph whose nodes are terms linked to one another by such properties as co-occurrence in text or grammatical morphology and learn the weights of their connections using graph search algorithms such as PageRank (Page et al. 1999). However, even these approaches fail to capture the deeper semantic meaning of documents. It is worth noting that, while semantic networks such as Google's Knowledge Graph bolster IR techniques with world facts and relationships, they are not concerned with a deeper analysis of query and document.

8 Conclusions and Future Work

We presented an investigation of an IR task in which sources containing sequences of events are matched to a query about the state of the world after those events. This task is challenging for traditional IR techniques, but key to simplifying access to information and reducing information overload.

We analyzed the problem from a commonsensical and intuitive perspective, and provided a formalization, based on action languages, of the desired reasoning. Although language $\mathcal{AL}$ is fundamental to our work, it is by itself insufficient, because it does not allow for the fine-grained reasoning needed for a clear determination of relevance in the presence of incomplete information and uncertainty. Thus, we presented an extension of $\mathcal{AL}$ suitable for our purpose. Finally, we defined an ASP-based procedure for automating the reasoning task.

In this paper, we have focused on introducing and studying the core IR task. Future work will address the connection with natural language processing algorithms and with available knowledge repositories, the development of an end-to-end system, and its quantitative evaluation.

References

BALDUCCINI, M. AND GELFOND, M. 2003. Diagnostic reasoning with A-Prolog. *Journal of Theory and Practice of Logic Programming (TPLP) 3,* 4–5 (Jul), 425–461.

BARAL, C. AND GELFOND, M. 2000. Reasoning Agents In Dynamic Domains. In *Workshop on Logic-Based Artificial Intelligence*. Kluwer Academic Publishers, 257–279.

BLANCO, R. AND LIOMA, C. 2012. Graph-Based Term Weighting for Information Retrieval. *Information Retrieval 15.1*, 54–92.

CAMPOS, R. 2015. Survey of Temporal Information Retrieval and Related Applications. *ACM Computing Surveys (CSUR) 47*, 2.

CARPINETO, C. AND RAMANO, G. 2012. A Survey of Automatic Query Expansion in Information Retrieval. *ACM Computing Surveys (CSUR) 44*, 1.

GELFOND, M. AND LIFSCHITZ, V. 1991. Classical Negation in Logic Programs and Disjunctive Databases. *New Generation Computing 9*, 365–385.

GLAVAS, G. AND SNAJDER, J. 2014. Event Graphs for Information Retrieval and Multi-Document Summarization. *Expert Systems with Applications 41,* 15, 6904–6916.

INCLEZAN, D. 2016. CoreALMlib: An ALM Library Translated from the Component Library. In *32nd International Conference on Logic Programming (ICLP16)*.

KORFHAGE, R. R. 1997. *Information Storage and Retrieval*. John Wiley and Sons, Inc.

LEBLANC, E. AND BALDUCCINI, M. 2016. Interpreting Natural Language Sources Using Transition Diagrams. In *Logic Programming with Constraints for Language Processing (CSLP2016)*, H. Christiansen and V. Dahl, Eds.

MANNING, C., CHRISTOPHER, D., RAGHAVAN, P., AND SCHÜTZE, H. 2008. *Introduction to Information Retrieval*. Vol. 1. Cambridge University Press.

MORALES, R., TU, P. H., AND SON, T. C. 2007. An Extension to Conformant Planning Using Logic Programming. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI'07)*, M. M. Veloso, Ed. 1991–1996.

NGUYEN, V., MITRA, A., AND BARAL, C. 2015. The NL2KR Platform for Building Natural Language Translation Systems. In *53rd Annual Meeting of the Association for Computational Linguistics (ACL-IJCNLP 2015)*. 899–908.

PAGE, L., BRIN, S., MOTWANI, R., AND WINOGRAD, T. 1999. The pagerank citation ranking: bringing order to the web.

SUCHANEK, F. M., KASNECI, G., AND WEIKUM, G. 2008. Yago: A Large Ontology from Wikipedia and WordNet. *Web Semantics: Science, Services and Agents on the World Wide Web 6,* 3, 203–217.