

Markov Logic Networks with Numerical Constraints

Melisachew Wudage Chekol and Jakob Huber and Christian Meilicke and Heiner Stuckenschmidt¹

Abstract. Markov logic networks (MLNs) have proven to be useful tools for reasoning about uncertainty in complex knowledge bases. In this paper, we extend MLNs with numerical constraints and present an efficient implementation in terms of a cutting plane method. This extension is useful for reasoning over uncertain temporal data. To show the applicability of this extension, we enrich log-linear description logics (DLs) with concrete domains (datatypes). Thereby, allowing to reason over weighted DLs with datatypes. Moreover, we use the resulting formalism to reason about temporal assertions in DBpedia, thus illustrating its practical use.

1 Motivation

Recent advances in data mining and information extraction have paved the way for the automatic construction of knowledge bases (KBs) from different sources, for instance, the NELL KB [26]. Often, the extraction tools used to construct such KBs produce weighted (or probabilistic) facts, due to the awareness that the implemented techniques cannot guarantee the completeness, correctness or consistency of the generated facts. Moreover, a large fraction of these facts can be temporal and may contain concrete data values, for example dates, times, latitudes/longitudes, numerical values measured in different units, and so on. Besides, the set of rules used to consolidate the KB will also be a set of weighted and unweighted rules with datatypes (numerical constraints). In this work, we propose an extension of Markov logic networks (MLNs) [32] that supports reasoning on numerical constraints.

We extend an MLNs inference approach (cutting plane [33]) to handle numerical constraints. This is difficult because numerical constraints may be infinite or continuous and reasoning about comparisons of two constraints is difficult to handle natively in MLNs. As an example, consider the following hard constraints which express: (1) a person whose age is between 13 and 19 is a teenager, and (2) a valid life span of a person is between 0 and 150 years:

$$\begin{aligned} 1 \forall x, y: & Person(x) \wedge age(x, y) \wedge y \geq 13 \wedge y \leq 19 \Rightarrow Teen(x) \\ 2 \forall p, bd, dd: & bdate(p, bd) \wedge ddate(p, dd) \Rightarrow lifeSpan(bd, dd), \\ & lifeSpan(bd, dd) = ((dd - bd) > 0) \wedge ((dd - bd) \leq 150) \end{aligned}$$

To the best of our knowledge, MLNs do not support numerical constraints (like the rules in the above example). Therefore, we propose an extended cutting plane approach to MLNs inference to introduce just the necessary constraint violations for the current MAP hypothesis taking also violated numerical constraints into account.

The proposed extension is useful for reasoning over uncertain temporal KBs as well as weighted Horn DLs with concrete domains. In this paper, we show that this extension is powerful enough to reason

over log-linear description logics (DLs) with concrete domains. Log-linear DLs [27] have been proposed as a way to combine tractable DLs with uncertainty reasoning while preserving the semantics of the DL part. They have also proven to be useful in situations where conflicting information from different sources has to be integrated into a coherent model. Successful applications include ontology matching [31] and information fusion [27]. So far the applicability of the combined model has been limited by the fact that it only supported reasoning on the schema level. In our work, we extend log-linear DLs to overcome its restriction to schema-level reasoning by introducing instances and numerical concrete domains (datatypes).

In DL, a concrete domain (CD) is a construct that can be used to define new classes by specifying constraints on attributes that have literal values (as opposed to relationships to other abstract entities). For instance, the axiom $Teenager \equiv Person \sqcap \geq_{13}(age) \sqcap \leq_{19}(age)$ defines a teenager as a person whose age is at least 13 and at most 19. While CDs, also referred to as *datatypes*, are a well-studied construct in classical DL (see for instance [22, 24]), this is barely the case in the probabilistic or log-linear extensions of DLs that have been proposed more recently (e.g. [7, 16, 11, 27, 21, 6, 35]). The contribution of this paper is the following:

- an extension of MLN with numerical constraints (MLN_{NC}),
- an extension of log-linear \mathcal{EL} with nominals and CDs by making use of MLN_{NC},
- an application of MLN_{NC} for reasoning about a specific sub-domain in the DBpedia knowledge base that contains numerical datatypes.
- an application of MLN_{NC} for debugging weighted temporal KBs.

Outline: In the next section we present the preliminaries: log-linear models, the lightweight ontology language \mathcal{EL}^{++} , and log-linear \mathcal{EL} obtained by combining log-linear models with \mathcal{EL}^{++} . In order to reason over uncertain temporal KBs, in Section 3, we extend MLNs with numerical constraints. In order to demonstrate additional benefits of this framework, we (1) extend log-linear \mathcal{EL} with CDs (in Section 4 and Section 5), and (2) present experimental results on debugging temporal KBs and computing the MAP state of an uncertain KB which contains datatypes (Section 6). Finally, we summarize the related work in Section 6, before making concluding remarks in Section 7.

2 Preliminaries

In this section, we present a brief summary of: Log-linear models, \mathcal{EL}^{++} , and log-linear \mathcal{EL} . For a detailed discussion on these subjects, we refer the reader to [2, 32, 34, 27] and the references therein.

¹ Data and Web Science Group, University of Mannheim, 68161 Mannheim, Germany, email: mel|jakob|christian|heiner@informatik.uni-mannheim.de

2.1 Log-Linear Models

A *log-linear model* is a compact representation of a probability distribution over assignments to a set of discrete-valued random variables $\mathcal{X} = \{X_1, \dots, X_n\}$ [17, 9]. The log-linear model is defined in terms of a set of feature functions $F = \{f_1(\mathbf{X}_1), \dots, f_k(\mathbf{X}_k)\}$, each of which is a function that defines a numerical value for each assignment \mathbf{x}_k to some subset $\mathbf{X}_k \subseteq \mathcal{X}$. Given a set of feature functions F , the parameters of the log-linear model are weights $W = \{w_k : f_k \in F\}$. The overall distribution is then defined as: $P(\mathbf{x}) = Z^{-1} \exp(\sum_{f_k \in F} w_k f_k(\mathbf{x}_k))$, where \mathbf{x}_k is the assignment to \mathbf{X}_k within \mathbf{x} , and Z is the normalization constant. A log-linear model induces a *Markov network* over \mathcal{X} , where there is an edge between every pair of variables $X_i, X_j \in \mathbf{X}_k$ that appear together in some feature $f_k(\mathbf{X}_k)$ and \mathbf{X}_k is a clique. Markov networks can also be encoded as a log-linear models by defining a feature function for every assignment of variables \mathbf{x}_c to a clique \mathbf{X}_c .

Markov Logic Networks (MLNs) can be seen as a first-order template language for log-linear models with binary variables. MLNs combine Markov networks and first-order logic (FOL) by attaching weights to first-order formulas and viewing these as templates for features of Markov networks [32]. An MNL L is a set of pairs (F_i, w_i) where F_i is a formula in FOL and w_i is a real number representing a weight. Together with a finite set of constants C , it defines a Markov Network $M_{L,C}$, where $M_{L,C}$ contains one node for each possible grounding of each predicate appearing in L . The value of the node is 1 if the ground predicate is true, and 0 otherwise. The probability distribution over possible worlds x specified by the ground Markov network $M_{L,C}$ is given by:

$$P(X = x) = \frac{1}{Z} \exp\left(\sum_{i=1}^F w_i n_i(x)\right)$$

where F is the number of formulas in the MLN and $n_i(x)$ is the number of true groundings of F_i in x . The groundings of a formula are formed simply by replacing its variables with constants in all possible ways. The *Herbrand base* HB is the set of all ground predicates (atoms) that can be constructed using the predicates in L and the constants in C . Each subset of the Herbrand base is a *Herbrand interpretation* specifying which ground atoms are true. A Herbrand interpretation H is a Herbrand model of L , written $\models_H L$, iff it satisfies all groundings of formulas in L .

There are two principal reasoning tasks in MLN, namely, MAP inference and Marginal inference. MAP inference is the task of finding the most probable world given some observations also referred to as evidence. Given the observed variables $E = e$, the MAP problem aims to find an assignment of all non-evidence (hidden) variables $X = x$ such that $\mathbf{I} = \operatorname{argmax}_x P(X = x \mid E = e)$. We denote by \mathbf{I} , the assignment x which leads P to be maximal, i.e., a MAP state. In order to compute a MAP state of an MLN, the problem can be formulated as an integer linear program (ILP) using the cutting plane inference algorithm.

2.2 \mathcal{EL}^{++}

\mathcal{EL}^{++} is the DL underlying the OWL 2 profile OWL-EL².

Syntax: Given a set of concept names \mathbf{N}_C , role names \mathbf{N}_R , individuals \mathbf{N}_I , and feature names \mathbf{N}_F , \mathcal{EL}^{++} concepts and roles are formed

according to the syntax given in Table 1. A concept in \mathcal{EL}^{++} is either a top, bottom concept, an atomic concept, a concrete concept or a complex concept (formed by conjunction and existential restriction). A *concrete domain* \mathbf{D} is a pair $(\Delta^{\mathbf{D}}, \mathbf{P}^{\mathbf{D}})$ with $\Delta^{\mathbf{D}}$ a set and $\mathbf{P}^{\mathbf{D}}$ a set of predicate names. Each $p \in \mathbf{P}$ is associated with an arity $n > 0$ and an extension $p^{\mathbf{D}} \subseteq (\Delta^{\mathbf{D}})^n$. The abstract and concrete domains are linked via a set of *feature names* \mathbf{N}_F . In Table 1, p denotes a predicate of some concrete domain \mathbf{D}_j and f_1, \dots, f_n are feature names. In this work, we consider only numerical CDs. However, our approach can easily be extended to handle other CDs. An \mathcal{EL}^{++} TBox contains a set of GCI (General Concept Inclusion) axioms, i.e., $C \sqsubseteq D$, as well as (RI) role inclusion axioms, i.e., $r_1 \circ r_2 \sqsubseteq r$.

Semantics: The semantics of \mathcal{EL}^{++} concepts and roles, shown in Table 1, is given by an interpretation function $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ which consists of a non-empty (abstract) domain $\Delta^{\mathcal{I}}$ and a mapping $\cdot^{\mathcal{I}}$ that assigns to each atomic concept $A \in \mathbf{N}_C$ a subset of $\Delta^{\mathcal{I}}$, to each abstract role $R \in \mathbf{N}_R$ a subset of $\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$, to each concrete relation $f \in \mathbf{N}_F$ a subset of $\Delta^{\mathcal{I}} \times \mathbf{D}$, and to each individual $a \in \mathbf{N}_I$ an element of $\Delta^{\mathcal{I}}$. The mapping $\cdot^{\mathcal{I}}$ is extended to all concepts and roles as shown in Table 1.

Table 1. The \mathcal{EL}^{++} with concrete domains.

Name	Syntax	Semantics
top	\top	$\Delta^{\mathcal{I}}$
bottom	\perp	\emptyset
nominal	$\{a\}$	$\{a^{\mathcal{I}}\}$
conjunction	$C \sqcap D$	$C^{\mathcal{I}} \cap D^{\mathcal{I}}$
existential restriction	$\exists r.C$	$\{x \in \Delta^{\mathcal{I}} \mid \exists y \in \Delta^{\mathcal{I}} : (x, y) \in r^{\mathcal{I}} \wedge y \in C^{\mathcal{I}}\}$
concrete domain	$p(f_1, \dots, f_n)$ for $p \in \mathbf{P}^{\mathbf{D}_j}$	$\{x \in \Delta^{\mathcal{I}} \mid \exists y_1, \dots, y_k \in \Delta^{\mathbf{D}_j} : f_i^{\mathcal{I}}(x) = y_i \text{ for } 1 \leq i \leq k \wedge (y_1, \dots, y_k) \in p^{\mathbf{D}_j}\}$
GCI	$C \sqsubseteq D$	$C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$
RI	$r_1 \circ r_2 \sqsubseteq r$	$r_1^{\mathcal{I}} \circ r_2^{\mathcal{I}} \subseteq r^{\mathcal{I}}$
concept assertion	$C(a)$	$a^{\mathcal{I}} \in C^{\mathcal{I}}$
role assertion	$r(a, b)$	$(a^{\mathcal{I}}, b^{\mathcal{I}}) \in r^{\mathcal{I}}$

Knowledge about specific objects can be expressed using concept and role assertions of the form $C(a)$ and $R(a, b)$. The axioms and assertions are contained in the TBox and ABox, respectively, which together form a knowledge base (KB). An \mathcal{EL}^{++} knowledge base $KB = (\mathcal{T}, \mathcal{A})$ consists of a set \mathcal{T} of general concept inclusion axioms (TBox) and role inclusion axioms, and possibly a set \mathcal{A} of assertional axioms (ABox). Despoina et al [8] have extended the inference/completion rules of \mathcal{EL}^{++} for the concrete domains by exploiting the notion of safety which keeps tractability of reasoning while enhancing expressivity. We will make use of these rules to provide datatype reasoning in log-linear \mathcal{EL} .

To simplify the translation of weighted \mathcal{EL}^{++} KBs into FOL, we first obtain the *normal* form of the KB in such a way that satisfiability is preserved [2, 18]. An \mathcal{EL}^{++} KB is in *normal* form $\tau(KB)$ if its axioms are in the following form:

$$C(a) \quad r(a, b) \quad A \sqsubseteq C \quad A \sqcap B \sqsubseteq C$$

² <http://www.w3.org/TR/owl2-profiles/>

$$\begin{aligned} \exists r.A \sqsubseteq C & \quad A \sqsubseteq \exists r.B & \quad r_1 \sqsubseteq r_2 \quad r_1 \circ r_2 \sqsubseteq r \\ A \sqsubseteq p(f_1, \dots, f_n) & \quad p(f_1, \dots, f_n) \sqsubseteq B \end{aligned}$$

where $A, B, C \in \mathcal{N}_C \cup \{\top\}$ and $C \in \{\perp\}$; $r, r_1, r_2 \in \mathcal{N}_R$; $f_1, \dots, f_n \in \mathcal{N}_F$; and $a, b, c \in \mathcal{N}_I$. For a finite set $\mathcal{N} \subseteq \mathcal{N}_C \cup \mathcal{N}_R$ of concept and role names the set of all normalized axioms constructible from \mathcal{N} is the union of (a) all normalized GCIs constructible from concept and role names in \mathcal{N} and the top and bottom concepts; and (b) all normalized RIs constructible from role names in \mathcal{N} .

Axioms of the form $C \sqsubseteq \{o\}$ are shown to be problematic for practical reasoning in the \mathcal{EL} family with nominals. However, experiments have shown that such axioms are rarely found in real world ontologies [14, 15]. In addition, by restricting their usage efficient practical reasoning can be achieved. Furthermore, for safe \mathcal{EL}^{++} , the inference rules given in [2] and proved in [15] are sound and complete. Therefore, for this study, we consider safe nominals as defined below.

Definition 1 (Kazakov et al [15]) An \mathcal{EL}^{++} concept C is safe if C has only occurrences of nominals in subconcepts of the form $\exists R.\{o\}$; C is negatively safe if C is either safe or a nominal. A GCI $C \sqsubseteq D$ is safe if C is negatively safe and D is safe. An \mathcal{EL}^{++} ontology is safe if all its concept inclusions are safe.

For instance, the axioms $\{a\} \sqsubseteq C$ and $\{a\} \sqsubseteq \exists r.\{b\}$ are safe. In this study, we consider safe \mathcal{EL}^{++} . It is possible to provide a probabilistic extension of safe \mathcal{EL}^{++} using MLNs. In which, a safe \mathcal{EL}^{++} KB can be seen as a set of hard constraints on the set of possible interpretations: if an interpretation violates even one axiom or assertion, it has zero probability. The basic idea in MLNs is to soften these constraints, i.e., when an interpretation violates one axiom or assertion in the KB it is less probable, but not impossible. The fewer axioms an interpretation violates, the more probable it becomes. Each axiom and assertion has an associated weight that reflects how strong a constraint is: the higher the weight, the greater the difference in log probability between an interpretation that satisfies the axiom and one that does not, other things being equal [32].

2.3 Log-Linear \mathcal{EL}

Log-linear \mathcal{EL} (LogEL) is an extension of \mathcal{EL}^{++} with log-linear models that enables reasoning over uncertain \mathcal{EL}^{++} TBoxes [27].

Syntax: A LogEL knowledge base $\text{KB} = (\text{KB}^D, \text{KB}^U)$ consists of a deterministic knowledge base KB^D and an uncertain knowledge base KB^U with $\text{KB}^D \not\models_{\mathcal{L}} \perp_{\mathcal{L}}$ and $\text{KB}^D \cap \text{KB}^U = \emptyset$ for some logic \mathcal{L} that supports a valid entailment relation $\models_{\mathcal{L}}$ and notion of contradiction $\perp_{\mathcal{L}}$. The uncertain KB is defined as $\text{KB}^U = \{\langle c_i, w_{c_i} \rangle\}$ where c_i is an axiom or assertion in \mathcal{L} and w_{c_i} is a real-valued weight assigned to c_i . The syntax of an axiom (resp. assertion) is similar to the underlying logic where an uncertain axiom (resp. assertion) has an associated weight as $\{\langle c_i, w_{c_i} \rangle\}$.

Semantics: The semantics of a LogEL knowledge base is based on joint probability distributions over the uncertain KB. Formally, for a given log-linear $\text{KB} = (\text{KB}^D, \text{KB}^U)$ and some KB' over the same signature, the probability of KB' is defined as:

$$P(\text{KB}') = \begin{cases} \frac{1}{Z} \exp \left(\sum_{\langle c_i, w_{c_i} \rangle \in \text{KB}^U: \text{KB}' \models c_i} w_{c_i} \right) & \text{if } \text{KB}' \not\models_{\mathcal{L}} \perp_{\mathcal{L}} \wedge \text{KB}' \models_{\mathcal{L}} \text{KB}^D, \\ 0 & \text{otherwise} \end{cases}$$

where Z is the normalization constant of the log-linear probability distribution P . Note that in MAP inference (i.e., obtaining the most probable KB) Z is not computed.

Example 1 Consider the following uncertain LogEL axioms: (1) a researcher is (probably) someone who published something, and (2) a famous researcher influenced (probably) someone who in turn has influenced someone else.

- (1) $\langle \text{Researcher} \sqsubseteq \exists \text{published}.\top, 0.8 \rangle$
- (2) $\langle \text{FamousResearcher} \sqsubseteq \exists \text{influenced}.\langle \exists \text{influenced}.\top \rangle \sqcap \text{Researcher}, 0.6 \rangle$

A LogEL KB can be normalized into an equivalent KB, thus, it can be mapped into first order predicates using a function φ as follows:

$$\begin{aligned} A \sqsubseteq C & \mapsto \text{sub}(A, C) & \exists r.A \sqsubseteq C & \mapsto \text{rsub}(A, r, C) \\ A \sqcap B \sqsubseteq C & \mapsto \text{int}(A, B, C) & r_1 \sqsubseteq r_2 & \mapsto \text{psub}(r_1, r_2) \\ A \sqsubseteq \exists r.B & \mapsto \text{rsup}(A, r, B) & r_1 \circ r_2 \sqsubseteq r & \mapsto \text{pcom}(r_1, r_2, r) \end{aligned}$$

The predicates in this listing are typed, r, r_1, r_2 are role names, A and B are concept names, and C is a concept name or the bottom concept (\perp). Note that φ is incomplete given an LogEL as defined above. In this work we will extend φ in order to deal with both ABox assertions and CDs (i.e., $p(f_1, \dots, f_n)$).

The translation of KB^D and KB^U results in unweighted (hard) and weighted first order formulas respectively. The hard formulas are used, together with the TBox completion rules $(F_1) - (F_9)$ to enforce $\text{KB}' \not\models_{\text{EL}} \perp_{\text{EL}}$ and $\text{KB}' \models_{\text{EL}} \text{KB}^D$ for any possible TBox KB' . Computing the MAP state of a LogEL knowledge base will always result in the most probable non contradictory subset of KB^U that entails the previously known axioms KB^D .

- (F_1) $\text{sub}(c, c)$
- (F_2) $\text{sub}(c, \top)$
- (F_3) $\text{sub}(c, c') \wedge \text{sub}(c', d) \Rightarrow \text{sub}(c, d)$
- (F_4) $\text{sub}(c, c_1) \wedge \text{sub}(c, c_2) \wedge \text{int}(c_1, c_2, d) \Rightarrow \text{sub}(c, d)$
- (F_5) $\text{sub}(c, c') \wedge \text{rsup}(c', r, d) \Rightarrow \text{rsup}(c, r, d)$
- (F_6) $\text{rsup}(c, r, d) \wedge \text{sub}(d, d') \wedge \text{rsub}(d', r, e) \Rightarrow \text{sub}(c, e)$
- (F_7) $\text{rsup}(c, r, d) \wedge \text{psub}(r, s) \Rightarrow \text{rsup}(c, s, d)$
- (F_8) $\text{rsup}(c, r_1, d) \wedge \text{rsup}(d, r_2, e) \wedge \text{pcom}(r_1, r_2, r_3) \Rightarrow \text{rsup}(c, r_3, e)$
- (F_9) $\neg \text{sub}(c, \perp)$

The above formulas (collectively denoted by \mathbf{F}) are universally quantified over all variables. They are partially derived from the completion rules of \mathcal{EL}^{++} [2]. \top and \perp are constant symbols representing the top and bottom concepts. Note that rule F_9 does not belong to the completion rules for \mathcal{EL}^{++} . This rule takes the notion of incoherence into account. An incoherent ontology is an ontology that contains an unsatisfiable concept, i.e., a concept that is subsumed by \perp . Usually, an unsatisfiable concept indicates that the ontology contains a contradictory set of axioms. An incoherent ontology is not necessarily inconsistent. Thus, we added rule F_9 which allows us to extend the notion of contradiction \perp_{EL} from inconsistency to incoherence. For more technical details on applying the principle of log-linear logic to OWL-EL, we refer the reader to [27]. We will extend these completion rules in order to deal with ABox assertions and CDs.

3 MLN with Numerical Constraints

We extend MLN with numerical constraints resulting in a formalism denoted MLN_{NC} . The constraints are predicates of the form $\theta \bowtie \psi$, where θ and ψ denote variables, numerical constants or algebraic expressions (that might contain elementary operators), and \bowtie is a binary operator which returns a truth value under some grounding.

Definition 2 (MLN with Numerical Constraints (MLN_{NC})) A numerical constraint NC is composed of numerical constants (such as elements of \mathbb{N} , \mathbb{I} , and so on), variables, elementary operators or functions (such as, $+$, $*$, $-$, \div , $\%$, $\sqrt{\quad}$), standard relations ($>$, $<$, $=$, \neq , \geq , \leq), and boolean operators (\wedge , \vee , \neg). An MLN_{NC} is a set of pairs (FC_i, w_i) where FC_i is a formula in FOL that may contain a NC and w_i is a real number representing the weight of formula FC_i .

Example 2 Using MLN_{NC} it is possible to represent the hard constraint: people born before 1850 are probably not alive: $\{\forall a, y : \text{person}(a) \wedge \text{born}(a, y) \wedge \text{NC}(y) \Rightarrow \text{dead}(y), \text{NC}(y) = y < 1850\}$.

MAP Inference in MLN_{NC} . A common inference task over MLNs is finding the most probable state of the world, i.e., finding a complete assignment to all ground atoms which maximizes the probability. A MAP query corresponds to an optimization problem with linear constraints and a linear objective function. Hence, it can be formulated and solved as an instance of an integer linear program (ILP) using the cutting plane approach, proposed in [34] and extended in [29]. Consequently, we extend the approach in [30] to compute a MAP state for a MLN_{NC} knowledge base. Cutting plane inference (CPI) operates between the grounding algorithm and the ILP solver. Instead of immediately adding one linear constraint for each ground clause to the ILP formulation, the ILP is initially formulated so as to enforce the evidence to hold in any solution. Based on the solution of this more compact ILP, one determines the violated constraints, adds these to the ILP, and solves the ILP again. This process is repeated until no constraints are violated by an intermediate solution. To elaborate, at the beginning of each CPI iteration it is necessary to determine the violated ground clauses G that are specified by the MLN and are in conflict with the intermediate solution. A binary ILP variable $x_\ell \in \{0, 1\}$ gets assigned to each grounded predicate occurring in a violated clause $g \in G$. The value of the variable x_ℓ is 1 if the respective literal ℓ is true and 0 if it is false. These variables are used to generate ILP constraints that are added to the ILP for each violated ground clause. For each clause $g \in G$, we define $L^+(g)$ as the set of ground atoms that occur unnegated in g and $L^-(g)$ as the set of ground atoms that occur negated in g . The transformation scheme depends on the weight $w_g \in \mathbb{R}$ of the violated clause g . It is also necessary to create a binary variable z_g for every g with $w_g \neq \infty$ that is used in the objective of the ILP. For every ground clause g with $w_g > 0$, the following constraint has to be added to the ILP.

$$\sum_{\ell \in L^+(g)} x_\ell + \sum_{\ell \in L^-(g)} (1 - x_\ell) \geq z_g$$

A ground atom ℓ that is set to false (or true if it appears negated) by evidence will not be included in the ILP as it cannot fulfill the respective constraint. For every g with weight $w_g < 0$, we add the following constraint to the ILP:

$$\sum_{\ell \in L^+(g)} x_\ell + \sum_{\ell \in L^-(g)} (1 - x_\ell) \leq (|L^+(g)| + |L^-(g)|)z_g$$

Algorithm 1 Extended cutting planes algorithm

Input: G : ground clauses; E : evidence clauses
Input: G_{NC} : ground clauses with NC
Input: $L_{\text{NC}}^+ \cup L_{\text{NC}}^-$: ground (positive or negated) numerical literals
Output: $H^{(t)}$: MAP state

- 1: **procedure** COMPUTEMAP(G, E)
- 2: $G_{\text{ILP}} \leftarrow E \cup \{\text{literal}(g) \in G \text{ and } w_g > 0\}$
- 3: Initial solution $H^{(0)} \leftarrow$ all atoms in G_{ILP}
- 4: $\text{ILP} \leftarrow \text{intoILP}(\forall g \in G_{\text{ILP}})$
- 5: $t \leftarrow 0$
- 6: **repeat**
- 7: $G_{\text{new}} \leftarrow \emptyset$
- 8: **for** $g \in G \setminus G_{\text{ILP}}$ **do**
- 9: **if** $((w_g > 0 \text{ or } w_g = \infty) \text{ and } H^{(t)} \not\models g)$
- 10: **or** $(w_g < 0 \text{ and } H^{(t)} \models g)$ **then**
- 11: **if** $g \in G_{\text{NC}}$ and g contains g_p and
- 12: $((\text{comp}(g_p) = 1 \text{ and } g_p \in L_{\text{NC}}^-) \text{ or}$
- 13: $(\text{comp}(g_p) = 0 \text{ and } g_p \in L_{\text{NC}}^+))$ **then**
- 14: $\text{add } g \setminus g_p \text{ to } G_{\text{new}}$
- 15: **else** $\text{add } g \text{ to } G_{\text{new}}$
- 16: **end if**
- 17: **end if**
- 18: **end for**
- 19: **if** $G_{\text{new}} \neq \emptyset$ **then**
- 20: $t \leftarrow t + 1$
- 21: $\text{ILP} \leftarrow \text{intoILP}(\forall g \in G_{\text{new}})$
- 22: $\text{add } g \in G_{\text{new}} \text{ to } G_{\text{ILP}}$
- 23: $H^{(t)} \leftarrow$ solution of ILP
- 24: **end if**
- 25: **until** $G_{\text{new}} = \emptyset$ **return** $H^{(t)}$
- 26: **end procedure**

The variable z_g expresses if a ground formula g is true considering the optimal solution of the ILP. However, for every g with weight $w_g = \infty$ this variable can be replaced with 1 as the respective formula cannot be violated in any solution:

$$\sum_{\ell \in L^+(g)} x_\ell + \sum_{\ell \in L^-(g)} (1 - x_\ell) \geq 1$$

Finally, the objective of the ILP sums up the weights of the (satisfied) ground formulas:

$$\max \sum_{g \in \mathcal{G}} w_g z_g.$$

The MAP problem can be solved as an ILP problem. Thus, the MAP state corresponds to the solution of the ILP in the last CPI iteration. It can be directly obtained from the solution as the assignment of the variables x_ℓ can be directly mapped to the optimal truth values for the ground predicates, i.e., $x_i = \text{true}$ if the corresponding ILP variable is 1 and $x_i = \text{false}$ otherwise.

We have extended the cutting planes algorithm proposed by Riedel [33] and optimized in [34] to the truth value of numerical predicates on-demand during each CPI iteration. This extension is sketched in Algorithm 1. The first initial MAP hypothesis is the set $H^{(0)}$. It contains all evidence clauses and ground clauses (containing only a single literal) with positive weights. These clauses have been added to the ILP and are denoted by G_{ILP} . At each CPI iteration, the violated constraints are identified and added to G_{new} . Thus, depending on the numerical predicates and the weight of the ground clauses, the violated constraints are determined (Lines 9–13). Consequently, for

every ground clause $g \in \mathbf{G} \setminus \mathbf{G}_{\text{ILP}}$, if g is not satisfied by the current hypothesis $H^{(t)}$ and $w_g > 0$ or $w_g = \infty$, or g is satisfied and $w_g < 0$, we test if g contains a NC. If g contains a NC (i.e., $g \in \mathbf{G}_{\text{NC}}$) and the numerical predicate g_p is in g , then we compute the numerical constraints corresponding to the predicate g_p with the current grounding, i.e., $\text{comp}(g_p)$. If $\text{comp}(g_p)$ is *true* and g_p is a negated literal $g_p \in L_{\text{NC}}^-$ or if $\text{comp}(g_p)$ is *false* and g_p is a positive literal $g_p \in L_{\text{NC}}^+$, we remove the numerical predicate from g and add g to \mathbf{G}_{new} . Note that we do not introduce ILP variables for numerical predicates as they will not be added to the \mathbf{G}_{new} , hence, to the ILP. Note also that our approach computes the truth value of numerical predicates on-demand, i.e., only when g is not satisfied by the current hypothesis. If g does not contain a numerical predicate, we add g to \mathbf{G}_{new} . If \mathbf{G}_{new} is empty, then the current hypothesis is optimal and we return it as MAP state. Otherwise, it iterates until an optimal solution is found.

Theorem 1 *Given a ground MLN_{NC} network, a set of observed variables E and a set of hidden variables X , an interpretation \mathbf{I} is a MAP state iff*

$$\mathbf{I} = \underset{X}{\text{argmax}} P(X | E) = \underset{X}{\text{argmax}} \sum_i w_i n_i(E, X).$$

Proof 1 (sketch) *It is proved that the MAP problem in MLN can be reduced to ILP using the CPI approach [34]. Likewise, we reduce MLN_{NC} 's MAP inference into a maximization problem in ILP using cutting planes. In Algorithm 1, if \mathbf{G} contains no ground numerical predicates, then the MAP inference in MLN_{NC} coincides with MAP inference in MLN. Otherwise, it is possible to compute a MAP state by transforming \mathbf{G} into ILP and integrating the truth value of the numerical predicates to the ILP by computing it outside the CPI setting as shown in Algorithm 1.*

4 Extending LogEL with Concrete Domains

We extend LogEL with uncertain knowledge expressed through numerical CDs.

Definition 3 *A numerical concrete domain \mathbf{D} is a pair $(\Delta^{\mathbf{D}}, \mathbf{P}^{\mathbf{D}})$ with $\Delta^{\mathbf{D}} \in \{\mathbb{N}, \mathbb{Z}, \mathbb{Q}, \mathbb{R}\}$ and $\mathbf{P}^{\mathbf{D}}$ a set of predicate names. Each $p \in \mathbf{P}$ is associated with an arity n and an extension $p^n \subseteq (\Delta^{\mathbf{D}})^n$.*

4.1 Unary concrete domains

We call a CD \mathbf{D} *unary* if all predicates $p \in \mathbf{P}^{\mathbf{D}}$ are unary, i.e., $n = 1$. In this section, we consider the unary predicates $\{<_v, \leq_v, >_v, \geq_v, =_v\}$ where $v \in \mathbb{R}$. The DL underlying LogEL admits safeness of CDs to retain tractability of reasoning. We present a unary CD \mathbf{R} inspired by safe numerical CDs introduced in [8]. *Safeness* is defined in terms of the position in which a predicate appears with respect to the subsumption (\sqsubseteq) relation in an axiom. If a predicate p appears on the right side of \sqsubseteq , then we call it a *positive predicate* ($p \in \mathbf{P}_+^{\mathbf{D}}$) and if on the left side, then it is a *negative predicate* ($p \in \mathbf{P}_-^{\mathbf{D}}$).

Definition 4 *A safe unary numerical CD \mathbf{R} is a triple $(\Delta^{\mathbf{R}}, \mathbf{P}_+^{\mathbf{R}}, \mathbf{P}_-^{\mathbf{R}})$ with $\Delta^{\mathbf{R}} \subseteq \mathbb{R}$ a set and $\mathbf{P}_+^{\mathbf{R}}, \mathbf{P}_-^{\mathbf{R}} \subseteq \{<_v, \leq_v, >_v, \geq_v, =_v\}$ are sets of positive and negative predicate names respectively.*

In the following, we consider an extension of LogEL with a safe unary numerical CD \mathbf{R} and refer to it as LogEL(\mathbf{R}).

Example 3 *We extend Example 1 with axioms that use safe numerical CDs where $<_{1850} \in \mathbf{P}_-^{\mathbf{D}}$ and $<_{1800} \in \mathbf{P}_+^{\mathbf{D}}$.*

$$\begin{aligned} & \langle \text{University} \sqcap <_{1850}(\text{foundedIn}) \\ & \quad \sqsubseteq \text{DistinguishedUniversity}, 0.4 \rangle \\ & \langle \text{OldUniversity} \sqsubseteq <_{1800}(\text{foundedIn}) \sqcap \text{University}, 0.6 \rangle \end{aligned}$$

Inference involving axioms that contain the CD \mathbf{R} can be done according to the following deduction/completion rules. In Definition 6, we will transform them into FOL formulas for reasoning over LogEL(\mathbf{R}) KB.

$$\frac{A \sqsubseteq B \quad B \sqsubseteq p_v(f)}{A \sqsubseteq p_v(f)} \quad \frac{A \sqsubseteq p_v(f) \quad p_{v'}(f) \sqsubseteq B \quad \text{eval}(p, v, v', v')}{A \sqsubseteq B}$$

where $p_v \in \mathbf{P}_+^{\mathbf{D}}, p_{v'} \in \mathbf{P}_-^{\mathbf{D}}$ and $v, v' \in \mathbb{R}$. In addition, the *eval* function checks if all possible values of the first *operator-value* pair (p, v) are covered by the possible values of the second *operator-value* pair (p', v') , when so, it evaluates to true otherwise false. The function *eval* is defined based on the domain \mathbb{R} and algebraic operators in \mathbf{R} . Some of the inequalities that are computed using the *eval* function are listed below:

$$\begin{aligned} \text{eval}(=, v, =, v') & := v = v' \\ \text{eval}(=, v, >, v') & := v > v' \end{aligned}$$

Example 4 *From the axioms of Example 3, it is possible to infer that $\text{OldUniversity} \sqsubseteq \text{DistinguishedUniversity}$ because $\text{eval}(<, 1800, <, 1850) = 1800 \leq 1850 = \text{true}$.*

In Section 5, we will see that the computation of the *eval* function integrates well with the proposed extension of the cutting planes algorithm.

4.2 Binary concrete domains

Unary CDs are not expressive enough to enable to describe the *relationship* between concrete values [22]. For instance, to describe a person that passed away at birth, we can use a binary CD as: $\text{DeadAtBirth} \sqsubseteq =(\text{birthDate}, \text{deathDate})$. Next, we introduce such expressive binary CDs. The main advantage of binary CDs over unary ones is that the former provides for a richer set of predicates [23]. Consequently, we introduce a CD \mathbf{T} which is based on real numbers and a set of unary and binary predicates. \mathbf{T} is defined as follows: $\mathbf{T} = (\Delta^{\mathbf{T}}, \mathbf{P}^{\mathbf{T}})$ where $\Delta^{\mathbf{T}} = \mathbb{R}$ and $\mathbf{P}^{\mathbf{T}}$ is defined as the smallest set containing the following predicates:

- unary predicates $\top_{\mathbf{T}}$ with $(\top_{\mathbf{T}})^{\mathbf{T}} = \mathbb{R}$ and $\perp_{\mathbf{T}}$ with $(\perp_{\mathbf{T}})^{\mathbf{T}} = \emptyset$,
- a unary predicate p_r for each $p \in \{<, \leq, =, \neq, \geq, >\}$ and each $r \in \mathbb{R}$ with $(p_r)^{\mathbf{T}} = \{r' \in \mathbb{R} \mid r' p r\}$, and
- binary predicates $p \in \{<, \leq, =, \neq, \geq, >\}$ with $(p)^{\mathbf{T}} = \{(r, r') \in \mathbb{R}^2 \mid r p r'\}$.

We will use the c \mathbf{T} for temporal reasoning over weighted LogEL(\mathbf{T}) KBs as shown in Section 6.2. In order to reason over such KBs, we design a set of domain specific rules which are used either for inferring a new knowledge or checking conflicts (aka. debugging). This is partly due to most of the current KBs do not contain TBox axioms that involve binary CDs. Note that it is possible to define a temporal CD based solely on Allen's interval relations which are often used in temporal reasoning [1]. Allen's relations describe the relationship between any two intervals over some temporal structure (for instance, a set of time points under a strict total ordering relation \prec) and can

be defined in terms of the interval endpoints. And thus, they can be expressed via the binary predicates of the CD \mathcal{T} . Hence, the CD \mathcal{T} is more expressive than a CD obtained by using Allen relations.

For the binary CDs we do not provide TBox completion rules, instead for the purpose of experimentation, we design a set of KB specific constraints. Examples of such rules are the following:

- (1) $\forall x, t1, t2 : Person(x) \wedge birthDate(x, t1) \wedge graduationDate(x, t2) \Rightarrow \langle t1, t2 \rangle$ 0.9
- (2) $\forall x, y, t1 : YoungAuthor(x) \wedge age(x, t1) \Rightarrow \langle t1, 20 \rangle$ 0.8

These rules are used in the experiments as discussed in Section 6.2.

5 Translating LogEL(R) and LogEL(T) into MLN_{NC}

As shown in previous work, the MAP state for a LogEL KB can be computed using a MLN that represents the first order transformation of the KB and entailment rules.

We have already presented the function φ for the translation of LogEL TBox axioms into FOL predicates. We now extend this function to map the CDs and ABox assertions of a LogEL KB into ground FOL predicates. Since the DL, \mathcal{EL}^{++} , underlying LogEL is equipped with nominals, ABox assertions can be converted into TBox axioms. Thus, with nominals, the ABox becomes syntactic sugar as shown below:

$$C(a) \Leftrightarrow \{a\} \sqsubseteq C, \quad r(a, b) \Leftrightarrow \{a\} \sqsubseteq \exists r.\{b\}$$

Instance checking in turn is directly reducible to subsumption checking in the presence of nominals. Thereby reducing ABox reasoning into TBox reasoning. Thus, in LogEL(R) and LogEL(T), we transform the ABox assertions into TBox axioms using nominals. We also extend the definition of normal forms for LogEL KBs with CDs as follows: for a LogEL KB = (KB^D, KB^U) , its normal form is given by: $\tau'(KB) = \tau(KB^D) \cup \bigcup_{(c_i, w_{c_i}) \in KB^U} \tau(c_i)$.

Definition 5 φ translates a normalized LogEL(R) and LogEL(T) KB's ABox and CDs into FOL formulas in MLN_{NC} as follows:

$$\begin{aligned} C(a) &\mapsto \text{sub}(\{a\}, C) \\ r(a, b) &\mapsto \text{rsup}(\{a\}, r, \{b\}) \\ A \sqsubseteq p_v(f) &\mapsto \text{sub}_+(A, p, v, f) \\ p_v(f) \sqsubseteq B &\mapsto \text{sub}_-(p, v, f, C) \end{aligned}$$

Definition 6 The translation of LogEL(R) TBox completion rules into FOL formulas \mathbf{F} in MLN_{NC} is given below:

$$\begin{aligned} (F_{10}) \quad &\text{sub}(a, b) \wedge \text{sub}_+(b, p, v, f) \Rightarrow \text{sub}_+(a, p, v, f) \\ (F_{11}) \quad &\text{sub}_+(a, p, v, f) \wedge \text{sub}_-(p', v', f, b) \wedge \\ &\text{eval}(p, v, p', v') \Rightarrow \text{sub}(a, b) \end{aligned}$$

Note that the FOL formulas \mathbf{F} corresponding to the translation of the completion rules are hard constraints. As the translation function φ is bijective, a possible world of \mathbf{F} and a given KB can be turned into a classified and consistent \mathcal{EL}^{++} (R) KB and vice versa as shown in the following lemma.

Lemma 1 Let KB be a normalized LogEL(R) (resp. LogEL(T)) knowledge base over a finite set \mathbf{N} of concept, role, and feature names and let HB be the Herbrand base of \mathbf{F} with respect to \mathbf{N} . If $KB' \subseteq KB$ is a classified and coherent knowledge base, then $\varphi(KB')$ is a Herbrand model of \mathbf{F} . If $HB' \subseteq HB$ is a Herbrand model of \mathbf{F} then $\varphi^{-1}(HB')$ is a classified and coherent knowledge base.

For a finite set of concept, role, individual and feature names \mathbf{F} , every normalized KB over the signature \mathbf{N} which is classified and coherent, corresponds to exactly one Herbrand model of \mathbf{F} .

MAP Inference Computing the MAP state of a given LogEL(R) (resp. or LogEL(T)) KB requires to translate the KB with the function φ (see Definition 5) to the equivalent Markov logic formalization. Then the completion rules of Definition 6 are added to this translation. The MAP state is computed with the help of Algorithm 1 applied to this input data. To do so, the evidence clauses $\varphi(KB)$ and the grounding of \mathbf{F} with respect to $\varphi(KB)$ are given as input to the algorithm. Besides, Algorithm 1 uses the external *comp* function whenever one of the grounding of the completion rules (F_{10}) and (F_{11}) requires to check a concrete grounding of the *eval* predicate. Note that the *comp* function can be easily implemented for the specified set of unary and binary numerical predicates (corresponding to a particular CD). Moreover, further extension of the proposed approach will also benefit from the fact that the externally defined *comp* function can be easily extended to deal with other predicates, for example, to support string comparisons. Applying the inverse translation function φ^{-1} to the MAP state, yields the most probable, classified, and consistent \mathcal{EL}^{++} (R) (resp. \mathcal{EL}^{++} (T)) KB.

Theorem 2 Given the following:

- a LogEL(R) (resp. LogEL(T)) normalized KB = (KB^D, KB^U) over a finite set \mathbf{N} of concept, role, individual and feature names,
- the Herbrand base HB of the formulas \mathbf{F} with respect to \mathbf{N} ,
- the set of ground formulas G_1 constructed from KB^D , and
- the set of ground formulas G_2 constructed from KB^U .

The most probable coherent and classified ontology is obtained with:

$$\varphi^{-1}(I) = \arg \max_{HB \supseteq I = G_1 \cup F} \left(\sum_{(g_j, w_j) \in G_2: I \neq g_j} w_j \right)$$

From Theorem 2 and the results in [27], the problem of computing the most probable, classified and coherent LogEL(R) (resp. LogEL(T)) KB is NP-hard.

6 Experiments

In the following we report about two types of experiments: temporal reasoning in LogEL(R) and LogEL(T) KBs.

6.1 Reasoning in LogEL(R)

In this experiment, we first illustrate in how far our formalism can be applied to a scenario where we reason with a knowledge base $KB = (KB^D, KB^U)$ with $KB^U = \emptyset$. In the second set of experiments we add weights to all assertions in KB^D , while adding at the same time a set of weighted erroneous assertions. We show that we achieve similar reasoning results in this noisy setting, while automatically detecting and removing erroneous numerical assertions.

6.1.1 Deterministic Temporal Reasoning

We apply our method to a specific subdomain of DBpedia that deals with researchers, their alma mater, their birth and death date, their publications together with the publishing date, and the influence relationships between researchers. For our experiments we worked with




Figure 1. A part of the elite universities from 1800 to 1925.

a slightly simplified signature. For example, we used the object property *influenced* for all assertions using one of the DBpedia properties *influenced*, *academic advisor*, and *notable student*. We extracted a dataset from DBpedia containing >11K persons, >1K universities, >10K books, and \approx 10K birth and death dates, >19K influenced assertions, and >700 publishing dates.

Then we added the axiom *FamousResearcher* $\equiv \exists \textit{influenced} . (\exists \textit{influenced} . \top)$ as well as a set of axioms for which the following axiom is an example.

$$\begin{aligned} & \exists \textit{hasStudent} . (\textit{FamousResearcher} \sqcap \\ & >_{1830}(\textit{born}) \sqcap \exists \textit{authorOf} . <_{1910}(\textit{published})) \\ & \sqsubseteq \textit{EliteUniversity}_{1850-1900} \end{aligned}$$

The first axiom defines a famous researcher as someone has influenced someone who in turn has influenced someone else. The second axiom states that an elite or distinguished university is one which had a student who was a famous researcher born after 1830 and published some work before 1910. With the help of these two axioms we can entail which academic institutions had a significant impact on the research community during the period 1850–1900. We created the same type of axioms for all 50 year periods starting in 1700 using a step size of 25 years. Then we computed the MAP state and retrieved the instances for all *EliteUniversity* concepts. A subset of the results of our query, with respect to the periods in the range from 1800 to 1925, is shown in Figure 1. Note that the number of *EliteUniversity* increases over time, which might also be caused by the fact that DBpedia is more complete with respect to more recent publications. Nevertheless, all of the universities from the results obtained are well known and distinguished universities.

6.1.2 Uncertain Temporal Reasoning

In this set of experiments we generated erroneous assertions specifying the publishing date of a book and the birth and death dates of persons. We injected a fraction of 20%, 40%, 60%, 80%, and finally 100% incorrect statements to the dataset. For instance, injecting 20% erroneous facts means that we added 20% additional wrong assertions for each of the three datatype properties. We randomly assigned weights in the range [0.5, 1.0] to the injected assertions and [0.8, 1.0] to each of the originally stated assertions.

We added a set of axioms which allow to detect that some combinations of birth, death and publishing dates are inconsistent. For instance, the following axiom expresses the fact that a person who is younger than 15 years of age cannot publish something.

$$\begin{aligned} & \textit{FamousResearcher} \equiv \exists \textit{influenced} . (\exists \textit{influenced} . \top) & 0.6 \\ & >_{1850}(\textit{born}) \sqcap \exists \textit{authorOf} . <_{1865}(\textit{published}) \sqsubseteq \perp & 0.82 \\ & \exists \textit{hasStudent} . (\textit{FamousResearcher} \sqcap \\ & >_{1830}(\textit{born}) \sqcap \exists \textit{authorOf} . <_{1910}(\textit{published})) \\ & \sqsubseteq \textit{EliteUniversity}_{1850-1900} & 0.75 \end{aligned}$$

We added similar axioms to express, e.g., that nobody can be born after he died. All these axioms use only unary predicates. Thus we had to generate the axioms for different time points. We did this with a step size of 25, i.e., for each type of axiom we generated a set of axioms to cover the time span from 1700 to 2000.

The MAP state for this input data will be a subset of the weighted input assertions, which might contain incorrect and correct assertions. The MAP state will also contain the instantiations of all *EliteUniversity* concepts. We compare the outcome of our approach against a gold standard, by assuming that each originally stated fact in DBpedia is correct, and each added fact is incorrect. Furthermore, we treat the *EliteUniversity* query result of the first experiment as a gold standard to measure in how far the results are negatively affected by the noisy setting.

Table 2. Precision (P) and recall (R) scores for computing the MAP state with an increasing number of injected erroneous assertions.

Injection		0%	20%	40%	60%	80%	100%
Query	P	1.00	0.92	0.81	0.76	0.66	0.61
	R	1.00	0.93	0.86	0.79	0.76	0.70
Repair	P	-	0.83	0.83	0.83	0.83	0.83
	R	-	0.62	0.62	0.61	0.61	0.60
Time (s)		390	631	1313	2351	3336	4599

The results of our experiments are shown in Table 2. Since we randomly assigned weights, we repeated each experiment 10 times and present average scores. We were able to compute meaningful results in highly inconsistent settings. Even in a setting where we added 100% incorrect assertions (as much correct assertions as incorrect assertions), we are still able to achieve a query precision of 61% and a recall of 70%. The good query results are caused by the relatively good results for the repair, where we computed the fraction of incorrect statements that have been removed (recall of repair) and the fraction of correct removal decisions (precision of repair). Both values are stable for different injection rates. The measured runtimes do not increase linearly with respect to the size of the input data. This is in line with our expectation, since the original data (0% injection) set will only require a materialization, while each added incorrect assertion might be involved in a conflict resulting finally in a non trivial optimization problem.

6.2 Debugging LogEL(T) KBs

The objective of this experiment is debugging temporal LogEL(T) KBs. We define a set of domain specific temporal rules (similar to those in the example of Section 4.2) based on a common sense understanding of this dataset. Typical examples of such rules are the following ones.

$$\begin{aligned} (3) \quad & \forall x, t1, t2 : \textit{person}(x) \wedge \textit{birthDate}(x, t1) \wedge \\ & \textit{deathDate}(x, t2) \Rightarrow <(t1, t2) & 0.7 \\ (4) \quad & \forall x, y, t1, t2 : \textit{almaMater}(x, y) \wedge \textit{established}(y, t1) \wedge \\ & \textit{deathDate}(x, t2) \Rightarrow <(t1, t2) & 0.5 \end{aligned}$$

Besides, we injected several types of erroneous facts to the extracted dataset. The results of our experiments are depicted in Table 3. The first column informs about the precision of the generated dataset. We have injected a fraction of 1%, 10%, 25%, 50%, 75%, and finally 100% incorrect statements to the dataset. Thus, the precision of the generated dataset varies from 0.99 (260K facts) to 0.5

(520K facts). For example, the last row of the table refers to a dataset where every second fact is incorrect. The recall of this dataset is always 1.0, because we never removed a fact of the extracted dataset. For all datasets, we have computed the MAP state, which corresponds to the repaired dataset. In doing so we computed precision, recall and F-measure for the debugging process and for the final outcome. The precision of the debugging process refers to the fraction of removed axioms or assertions that were indeed incorrect; recall refers in this context to the fraction of incorrect axioms or assertions that have been removed. The precision and recall of the repaired dataset is computed by comparing it to the originally extracted dataset. The rightmost column (ΔF) informs about the gain in F-measure that we computed by comparing the F-measure of the input dataset with the F-measure of the repaired dataset.

Table 3. Precision (P), Recall (R) and F-measure (F) for debugging process and repaired dataset.

Input	Debugging			Repaired Dataset			ΔF
	P	R	F	P	R	F	
0.99	0.80	0.63	0.70	1.00	1.00	1.00	0.002
0.91	0.80	0.64	0.71	0.97	0.98	0.97	0.022
0.80	0.81	0.65	0.72	0.92	0.96	0.94	0.050
0.67	0.82	0.65	0.73	0.84	0.93	0.88	0.084
0.57	0.83	0.65	0.73	0.78	0.90	0.83	0.106
0.50	0.84	0.65	0.73	0.72	0.87	0.79	0.119

The results show that we are able to improve the data quality of an erroneous dataset. This is depicted in the increase of F-measure in the ΔF column for all test cases. With respect to the test case where every second fact is incorrect, we are able to increase the precision by 0.22 (from 0.5 up to 0.72), while recall is only reduced to 0.87 (from 1.0). Most of the removed facts are indeed incorrect (a repair precision of > 0.8 means that at least for 4 of 5 removals are proper removals), while we are able to detect more than half of the incorrect facts (debugging recall of the repair is > 0.5).

7 Related Work

Since in MLNs all variables and features are discrete, Hybrid MLNs that allow for both discrete and continuous variables in addition to limited numerical constraints have been introduced [36]. While our approach does not support continuous variables, the numerical constraints that we consider here are more expressive and the ability to compute them outside the cutting planes algorithm, enables us for far more richer constraints and functions. In the context of probabilistic programming, supporting numerical constraints in probabilistic reasoning is also explored [5, 4, 3].

The study of extending DLs to handle uncertainty and vagueness has gained momentum recently. There have been several proposals to add probabilities to various DLs. Probabilistic DLs can be classified in several dimensions. One possible classification is on the reasoning mechanism used: Markov logic networks (MLNs) and Bayesian networks. There exist some studies that employ MLNs to extend various DLs. The study in [21] extends \mathcal{EL}^{++} with probabilistic uncertainty based on the annotation of axioms using MLNs. The main focus of this work is ranking queries in descending order of probability of atomic inferences which is different from the objective of this paper. Another study in [27], presents a probabilistic extension of the DL \mathcal{EL}^{++} without nominals and CDs in MLN in order to find the most probable coherent ontology. In doing so, they have developed a

reasoner for probabilistic OWL-EL called ELOG [29]. In this study, we extend this work in order to deal with CDs in addition to nominals and instances. In databases, MLNs have been used to create a probabilistic datalog called Datalog+/. It is an extension of datalog that allows to express ontological axioms by using rule-based constraints [10]. The probabilistic extension of Datalog+/- uses MLNs as the underlying probabilistic semantics. The focus of this work is on scalable threshold query answering which is different from that of this work.

Other literatures extend DLs with Bayesian networks. Some notable works include: an extension of \mathcal{EL} with Bayesian networks called \mathcal{BEL} is presented in [6]. They study the complexity of reasoning under \mathcal{BEL} to show that reasoning is intractable. However, their work does not discuss probabilities in the ABox and concrete domains are excluded. On the other hand, in [7], they added uncertainty to DL-Lite based on Bayesian networks. Additionally, they have shown that satisfiability test and query answering in probabilistic DL-Lite can be reduced to satisfiability test and query answering in the DL-Lite family. Further, it is proved that satisfiability checking and union of conjunctive query answering can be done in LogSpace in the data complexity. Query answering in probabilistic OWL QL, where ABox assertions are probabilistic and TBox axioms are deterministic, has been studied in [13]. They prove that only very simple conjunctive queries can be answered in PTime, while most queries are #P-hard.

It is possible to extend DLs with uncertainty by using Halpern's probabilistic first order logic [12]. These studies include Prob- \mathcal{ALC} [25], P- $\mathcal{SHOIN}(\mathbf{D})$ [35]. On the other hand, P- $\mathcal{SHIQ}(\mathbf{D})$ uses probabilistic lexicographic entailment from probabilistic default reasoning [19]. P- $\mathcal{SHOIN}(\mathbf{D})$ and P- $\mathcal{SHIQ}(\mathbf{D})$ support datatype reasoning in a probabilistic setting. While these logics are different from the one studied in this paper, they also do not support expressive CDs. In addition, while LogEL is based on log-linear models, P- $\mathcal{SHIQ}(\mathbf{D})$ is based on Nilsson's probabilistic logic [28] and P- $\mathcal{SHOIN}(\mathbf{D})$ is based on Halpern's probabilistic FOL; both of these formalisms are different from log-linear models. A survey of probabilistic extensions of DLs can be found in [20].

Consequently, as discussed above, most of the studies that involve extending DLs to deal with uncertainty by using either Bayesian or MLNs often excluded CDs. This is partly due to either the lack of supporting features or the difficulty in dealing with them. In this paper, we studied a novel way of dealing with uncertainty involving CDs by log-linear models extended with numerical constraints.

8 Conclusion

In this paper, we extended MLN with numerical constraints. To show its usefulness, we have applied this framework to extend LogEL with instances and CDs. This work can be seen as a first step towards more practical uses of log-linear models with numerical constraints. We have illustrated the practical merits of the approach using an example application. What needs to be done is a systematic evaluation of the approach. Furthermore, there are two directions for improving the formalism. On the one hand, we can further extend the modeling capabilities by including other kinds of datatypes and probability distributions over attribute values. On the other hand, we will investigate tractable subsets of extended MLNs to enable polynomial time reasoning.

REFERENCES

- [1] James F Allen, ‘Maintaining knowledge about temporal intervals’, *Communications of the ACM*, **26**(11), 832–843, (1983).
- [2] Franz Baader, Sebastian Brandt, and Carsten Lutz, ‘Pushing the envelope’, in *IJCAI*, volume 5, pp. 364–369, (2005).
- [3] Vaishak Belle, Andrea Passerini, and Guy Van den Broeck, ‘Probabilistic inference in hybrid domains by weighted model integration’, in *Proceedings of 24th International Joint Conference on Artificial Intelligence (IJCAI)*, (2015).
- [4] Vaishak Belle, Guy Van den Broeck, and Andrea Passerini, ‘Hashing-based approximate probabilistic inference in hybrid domains’, in *Proceedings of the 31st Conference on Uncertainty in Artificial Intelligence (UAI)*, (2015).
- [5] Rodrigo de Salvo Braz, Ciaran O’Reilly, Vibhav Gogate, and Rina Dechter, ‘Probabilistic inference modulo theories’, (2016). to appear.
- [6] Ismail Ilkan Ceylan and Rafael Penaloza, ‘Bayesian description logics’, in *Proceedings of the 27th International Workshop on Description Logics (DL 2014). CEUR Workshop Proceedings*, volume 1193, pp. 447–458, (2014).
- [7] Claudia d’Amato, Nicola Fanizzi, and Thomas Lukasiewicz, ‘Tractable reasoning with bayesian description logics’, in *Scalable Uncertainty Management*, 146–159, Springer, (2008).
- [8] M Despoina, Y Kazakov, and I Horrocks, ‘Tractable extensions of the description logic el with numerical datatypes’, *Journal of Automated Reasoning*, (2011).
- [9] Lise Getoor, *Introduction to statistical relational learning*, MIT press, 2007.
- [10] Georg Gottlob, Thomas Lukasiewicz, Maria Vanina Martinez, and Gerardo I Simari, ‘Query answering under probabilistic uncertainty in datalog+/- ontologies’, *Annals of Mathematics and Artificial Intelligence*, **69**(1), 37–72, (2013).
- [11] Víctor Gutiérrez-Basulto, Jean Christoph Jung, Carsten Lutz, and Lutz Schröder, ‘A closer look at the probabilistic description logic prob-el’, in *AAAI*, (2011).
- [12] Joseph Y Halpern, ‘An analysis of first-order logics of probability’, *Artificial intelligence*, **46**(3), 311–350, (1990).
- [13] Jean Christoph Jung and Carsten Lutz, ‘Ontology-based access to probabilistic data with owl ql’, in *The Semantic Web–ISWC 2012*, 182–197, Springer, (2012).
- [14] Yevgeny Kazakov, Markus Krötzsch, and František Simančík, ‘Practical reasoning with nominals in the \mathcal{EL} family of description logics’, in *Proceedings of the 13th International Conference on Principles of Knowledge Representation and Reasoning (KR’12)*, eds., Gerhard Brewka, Thomas Eiter, and Sheila A. McIlraith, pp. 264–274. AAAI Press, (2012).
- [15] Yevgeny Kazakov, Markus Krötzsch, and František Simančík, ‘The incredible ELK: From polynomial procedures to efficient reasoning with \mathcal{EL} ontologies’, *Journal of Automated Reasoning*, **53**, 1–61, (2013).
- [16] Pavel Klinov, ‘Pronto: a non-monotonic probabilistic description logic reasoner’, in *Proceedings of the 5th European semantic web conference on The semantic web: research and applications*, pp. 822–826. Springer-Verlag, (2008).
- [17] Daphne Koller and Nir Friedman, *Probabilistic graphical models: principles and techniques*, MIT press, 2009.
- [18] Markus Krötzsch, ‘Efficient rule-based inferencing for owl el’, in *IJCAI*, volume 11, pp. 2668–2673, (2011).
- [19] Thomas Lukasiewicz, ‘Expressive probabilistic description logics’, *Artificial Intelligence*, **172**(6), 852–883, (2008).
- [20] Thomas Lukasiewicz, ‘Expressive probabilistic description logics’, *Artificial Intelligence*, **172**(6), 852–883, (2008).
- [21] Thomas Lukasiewicz, Maria Vanina Martinez, Giorgio Orsi, and Gerardo I. Simari, ‘Heuristic ranking in tightly coupled probabilistic description logics’, in *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence, Catalina Island, CA, USA, August 14-18, 2012*, pp. 554–563, (2012).
- [22] Carsten Lutz, ‘Description logics with concrete domains—a survey’, in *Advances in Modal Logic 4, papers from the fourth conference on “Advances in Modal logic,” held in Toulouse (France) in October 2002*, pp. 265–296, (2002).
- [23] Carsten Lutz, ‘Combining interval-based temporal reasoning with general TBoxes’, *Artificial Intelligence*, **152**(2), 235–274, (2004).
- [24] Carsten Lutz and Maja Miličić, ‘A tableau algorithm for description logics with concrete domains and general tboxes’, *Journal of Automated Reasoning*, **38**(1-3), 227–259, (2007).
- [25] Carsten Lutz and Lutz Schröder, ‘Probabilistic description logics for subjective uncertainty’, in *KR*, (2010).
- [26] T. Mitchell, W. Cohen, E. Hruschka, P. Talukdar, J. Betteridge, A. Carlson, B. Dalvi, M. Gardner, B. Kisiel, J. Krishnamurthy, N. Lao, K. Mazaitis, T. Mohamed, N. Nakashole, E. Platanios, A. Ritter, M. Samadi, B. Settles, R. Wang, D. Wijaya, A. Gupta, X. Chen, A. Saparov, M. Greaves, and J. Welling, ‘Never-ending learning’, in *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-15)*, (2015).
- [27] Mathias Niepert, Jan Noessner, and Heiner Stuckenschmidt, ‘Log-linear description logics’, in *IJCAI*, pp. 2153–2158, (2011).
- [28] Nils J Nilsson, ‘Probabilistic logic’, *Artificial intelligence*, **28**(1), 71–87, (1986).
- [29] Jan Noessner and Mathias Niepert, ‘Elog: a probabilistic reasoner for owl el’, in *Web Reasoning and Rule Systems*, 281–286, Springer, (2011).
- [30] Jan Noessner, Mathias Niepert, and Heiner Stuckenschmidt, ‘Rockit: Exploiting parallelism and symmetry for map inference in statistical relational models’, in *Twenty-Seventh AAAI Conference on Artificial Intelligence*, (2013).
- [31] Jan Noessner, Heiner Stuckenschmidt, Christian Meilicke, and Mathias Niepert, ‘Completeness and optimality in ontology alignment debugging’, in *Proceedings of the 9th International Workshop on Ontology Matching collocated with the 13th International Semantic Web Conference (ISWC 2014)*, pp. 25–36, Riva del Garda, Trentino, Italy, (2014).
- [32] Matthew Richardson and Pedro Domingos, ‘Markov logic networks’, *Machine learning*, **62**(1-2), 107–136, (2006).
- [33] S Riedel, ‘Improving the accuracy and efficiency of map inference for markov logic’, in *Proceedings of the 24th Conference on Uncertainty in Artificial Intelligence, UAI 2008*, pp. 468–475, (2008).
- [34] Sebastian Riedel, ‘Improving the accuracy and efficiency of map inference for markov logic’, *arXiv preprint arXiv:1206.3282*, (2012).
- [35] Fabrizio Riguzzi, Elena Bellodi, Evelina Lamma, and Riccardo Zese, ‘Reasoning with probabilistic ontologies’, in *Proceedings of the 24th International Joint Conference on Artificial Intelligence*, pp. 4310–4316. AAAI Press, (2015).
- [36] Jue Wang and Pedro M Domingos, ‘Hybrid markov logic networks.’, in *AAAI*, volume 8, pp. 1106–1111, (2008).