

# Towards a Toolbox for Relational Probabilistic Knowledge Representation, Reasoning, and Learning

Marc Finthammer<sup>1</sup>, Sebastian Loh<sup>2</sup>, and Matthias Thimm<sup>2</sup>

<sup>1</sup> Department of Computer Science, FernUniversität in Hagen, Germany

<sup>2</sup> Department of Computer Science, Technische Universität Dortmund, Germany

**Abstract.** This paper presents KREATOR, a versatile and easy-to-use toolbox for statistical relational learning currently under development. The research on combining probabilistic models and first-order theory put forth a lot of different approaches in the past few years. While every approach has advantages and disadvantages the variety of prototypical implementations make thorough comparisons of different approaches difficult. KREATOR aims at providing a common and simple interface for representing, reasoning, and learning with different relational probabilistic approaches. We give an overview on the system architecture of KREATOR and illustrate its usage.

## 1 Introduction

*Probabilistic inductive logic programming* (or *statistical relational learning*) is a very active field in research at the intersection of logic, probability theory, and machine learning, see [1, 2] for some excellent overviews. This area investigates methods for representing probabilistic information in a relational context for both reasoning and learning. Many researchers developed liftings of propositional probabilistic models to the first-order case in order to take advantage of methods and algorithms already developed. Among these are the well-known Bayesian logic programs [3] and Markov logic networks [4] which extend respectively Bayes nets and Markov nets [5] and are based on knowledge-based model construction [6]. Other approaches also employ Bayes nets for their theoretical foundation like logical Bayesian networks [7] and relational Bayesian networks [8]; or they are influenced by other fields of research like probabilistic relational models [9] by database theory and P-log [10] by answer set programming. There are also some few approaches to apply maximum entropy methods to the relational case [11, 12]. But because of this variety of approaches and the absence of a common interface there are only few comparisons of different approaches, see for example [13, 14].

In this paper we describe the KREATOR toolbox, a versatile integrated development environment for knowledge engineering in the field of statistical relational learning. KREATOR is currently under development and part of the

ongoing KREATE project<sup>3</sup> which aims at developing a common methodology for learning, modelling and inference in a relational probabilistic framework. As statistical relational learning is a (relatively) young research area there are many different proposals for integrating probability theory in first-order logic, some of them mentioned above. Although many researchers have implementations of their approaches available, most of these implementations are prototypical, and in order to compare different approaches one has to learn the usage of different tools. KREATOR aims at providing a common interface for different approaches to statistical relational learning and to support the researcher and knowledge engineer in developing knowledge bases and using them in a common and easy-to-use fashion. Currently, the development of KREATOR is still in a very early stage but already supports Bayesian logic programs, Markov logic networks, and in particular a new approach for using maximum entropy methods in a relational context [11].

The rest of this paper is organized as follows. In Sec. 2 we give an overview on the approaches of statistical relational learning that are currently supported by KREATOR, i. e. Bayesian logic programs, Markov logic networks, and the relational maximum entropy approach. We go on in Sec. 3 with presenting the system architecture of KREATOR and motivate the main design choices. In Sec. 4 we give a short manual-style overview on the usage of KREATOR and in Sec. 5 we give some hints on future work and conclude.

## 2 Relational Probabilistic Knowledge Representation

In the following we give some brief overview on frameworks for relational probabilistic reasoning that are already implemented in KREATOR. These are Bayesian logic programs originally due to Kersting et. al. [3], Markov logic networks originally due to Domingos et. al. [4], and a framework employing reasoning with maximum entropy methods that is currently in development [11]. We illustrate the use of these frameworks on a common example, the well-known burglary example [5, 1].

*Example 1.* We consider a scenario where someone—let’s call him *James*—is on the road and gets a call from his neighbor saying that the alarm of James’ house is ringing. James has some uncertain beliefs about the relationships between burglaries, types of neighborhoods, natural disasters, and alarms. For example, he knows that if there is a tornado warning for his home place, then the probability of a tornado triggering the alarm of his house is 0.9. A reasonable information to infer from his beliefs and the given information is “What is the probability of an actual burglary?”.

### 2.1 Bayesian Logic Programs

Bayesian logic programming is an approach to combine logic programming and Bayesian networks [3]. Bayesian logic programs (BLPs) use a standard logic

<sup>3</sup> <http://www.fernuni-hagen.de/wbs/research/kreate/index.html>

programming language and attach to each logical clause a set of probabilities, which define a conditional probability distribution of the head of the clause given specific instantiations of the body of the clause.

In contrast to first-order logic, BLPs employ an extended form of predicates and atoms. In BLPs, *Bayesian predicates* are predicates that feature an arbitrary set as possible states that are not necessarily the boolean values  $\{\text{true}, \text{false}\}$ . For example, the Bayesian predicate *bloodtype/1* may represent the blood type of a person using the possible states  $S(\text{bloodtype}) = \{a, b, ab, 0\}$  [3]. Analogously to first-order logic, Bayesian predicates can be instantiated to *Bayesian atoms* using constants and variables and then each ground Bayesian atom represents a single random variable. If  $A$  is a Bayesian atom of the Bayesian predicate  $p$  we set  $S(A) = S(p)$ .

**Definition 1 (Bayesian Clause, Conditional Probability Distribution).**

A Bayesian clause  $c$  is an expression  $(H \mid B_1, \dots, B_n)$  with Bayesian atoms  $H, B_1, \dots, B_n$ . With a Bayesian clause  $c$  with the form  $(H \mid B_1, \dots, B_n)$  we associate a function  $\text{cpd}_c : S(H) \times S(B_1) \times \dots \times S(B_n) \rightarrow [0, 1]$  that fulfills

$$\forall b_1 \in S(B_1), \dots, b_n \in S(B_n) : \sum_{h \in S(H)} \text{cpd}_c(h, b_1, \dots, b_n) = 1 \quad .$$

We call  $\text{cpd}_c$  a conditional probability distribution. Let  $\text{CPD}_p$  denote the set of all conditional probability distributions  $\{\text{cpd}_{H \mid B_1, \dots, B_n} \mid H \text{ is an atom of } p\}$ .

A function  $\text{cpd}_c$  for a Bayesian clause  $c$  expresses the conditional probability distribution  $P(\text{head}(c) \mid \text{body}(c))$  and thus partially describes an underlying probability distribution  $P$ .

*Example 2.* We represent Ex. 1 as a set  $\{c_1, c_2, c_3\}$  of Bayesian clauses with

$$\begin{aligned} c_1 : & \quad (\text{alarm}(\text{X}) \mid \text{burglary}(\text{X})) \\ c_2 : & \quad (\text{alarm}(\text{X}) \mid \text{lives\_in}(\text{X}, \text{Y}), \text{tornado}(\text{Y})) \\ c_3 : & \quad (\text{burglary}(\text{X}) \mid \text{neighborhood}(\text{X})) \end{aligned}$$

where  $S(\text{tornado}/1) = S(\text{lives\_in}/2) = S(\text{alarm}) = S(\text{burglary}) = \{\text{true}, \text{false}\}$  and  $S(\text{neighborhood}) = \{\text{good}, \text{average}, \text{bad}\}$ . For each Bayesian clause  $c_i$ , we define a function  $\text{cpd}_{c_i}$  which expresses our subjective beliefs, e. g., for clause  $c_2$  we define

$$\begin{aligned} \text{cpd}_{c_2}(\text{true}, \text{true}, \text{true}) &= 0.9 & \text{cpd}_{c_2}(\text{true}, \text{true}, \text{false}) &= 0.01 \\ \text{cpd}_{c_2}(\text{true}, \text{false}, \text{true}) &= 0 & \text{cpd}_{c_2}(\text{true}, \text{false}, \text{false}) &= 0 \\ \text{cpd}_{c_2}(\text{false}, \text{true}, \text{true}) &= 0.1 & \text{cpd}_{c_2}(\text{false}, \text{true}, \text{false}) &= 0.99 \\ \text{cpd}_{c_2}(\text{false}, \text{false}, \text{true}) &= 1 & \text{cpd}_{c_2}(\text{false}, \text{false}, \text{false}) &= 1 \end{aligned}$$

Considering clauses  $c_1$  and  $c_2$  in Ex. 2 one can see that it is possible to have multiple clauses with the same head. BLPs facilitate *combining rules* in order to aggregate probabilities that arise from applications of different Bayesian clauses. A combining rule  $\text{cr}_p$  for a Bayesian predicate  $p/n$  is a function  $\text{cr}_p :$

$\mathfrak{P}(\text{CPD}_p) \rightarrow \text{CPD}_p$  that assigns to the conditional probability distributions of a set of Bayesian clauses a new conditional probability distribution that represents the *joint* probability distribution obtained from aggregating the given clauses<sup>4</sup>. For example, given clauses  $c_1 = (b(X) \mid a_1(X))$  and  $c_2 = (b(X) \mid a_2(X))$  the result  $f = \text{cr}_b(\{\text{cpd}_{c_1}, \text{cpd}_{c_2}\})$  of the combining rule  $\text{cr}_b$  is a function  $f : S(b) \times S(a_1) \times S(a_2) \rightarrow [0, 1]$ . Appropriate choices for such functions are *average* or *noisy-or*, cf. [3].

*Example 3.* We continue Ex. 2. Suppose *noisy-or* to be the combining rule for *alarm*. Then the joint conditional probability distribution  $\text{cpd}_{c'}$  for  $c' = (\text{alarm}(X) \mid \text{burglary}(X), \text{lives\_in}(X, Y), \text{tornado}(Y))$  can be computed via

$$\text{cpd}_{c'}(t_1, t_2, t_3, t_4) = Z * (1 - (1 - \text{cpd}_{c_1}(t_1, t_2)) * (1 - \text{cpd}_{c_2}(t_1, t_3, t_4)))$$

for any  $t_1, t_2, t_3, t_4 \in \{\text{true}, \text{false}\}$ . Here,  $Z$  is a normalizing constant for maintaining the property of conditional probability distributions to sum up to one for any specific head value.

Now we are able to define Bayesian logic programs as follows.

**Definition 2 (Bayesian Logic Program).** A Bayesian logic program  $B$  is a tuple  $B = (C, D, R)$  with a (finite) set of Bayesian clauses  $C = \{c_1, \dots, c_n\}$ , a set of conditional probability distributions (one for each clause in  $C$ )  $D = \{\text{cpd}_{c_1}, \dots, \text{cpd}_{c_n}\}$ , and a set of combining functions (one for each Bayesian predicate appearing in  $C$ )  $R = \{\text{cr}_{p_1}, \dots, \text{cr}_{p_m}\}$ .

Semantics are given to Bayesian logic programs via transformation into the propositional case, i. e. into Bayesian networks [5]. Given a specific (finite) universe  $U$  a Bayesian network  $BN$  can be constructed by introducing a node for every grounded Bayesian atom in  $B$ . Using the conditional probability distributions of the grounded clauses and the combining rules of  $B$  a (joint) conditional probability distribution can be specified for any node in  $BN$ . If  $BN$  is acyclic this transformation uniquely determines a probability distribution  $P$  on the grounded Bayesian atoms of  $B$  which can be used to answer queries.

A detailed description of the above (declarative) semantics and an equivalent procedural semantics which is based on SLD resolution are given in [3].

## 2.2 Markov Logic Networks

Markov logic [4] establishes a framework which combines Markov networks [5] with first-order logic to handle a broad area of statistical relational learning tasks. The Markov logic syntax complies with first-order logic<sup>5</sup>, however each formula is quantified by an additional weight value. The semantics of a set of Markov logic formulas is explained by a probability distribution over possible worlds. A possible world assigns a truth value to every possible ground atom

<sup>4</sup>  $\mathfrak{P}(S)$  denotes the power set of a set  $S$ .

<sup>5</sup> Although Markov logic also covers functions, we will omit this fact, and only consider constants.

(constructible from the set of predicates and the set of constants). The probability distribution is calculated as a log-linear model over weighted ground formulas.

The fundamental idea in Markov logic is that first-order formulas are not handled as hard constraints. Instead, each formula is more or less softened depending on its weight. So a possible world *may* violate a formula without necessarily receiving a zero probability. Rather a world is more probable, the less formulas it violates. A formula's weight specifies how strong the formula is, i. e. how much the formula influences the probability of a satisfying world versus a violating world. This way, the weights of all formulas influence the determination of a possible world's probability in a complex manner. One clear advantage of this approach is that Markov logic can directly handle contradictions in a knowledge base, since the (contradictory) formulas are weighted against each other anyway. Furthermore, by assigning appropriately high weight values to certain formulas, it can be enforced that these formulas will be handled as hard constraints, i. e. any world violating such a hard formula will have a zero probability. Thus, Markov logic also allows the processing of purely logical first-order formulas.

**Definition 3 (Markov logic network).** A Markov logic network (MLN)  $L$  is a set of first-order logic formulas  $F_i$ , where each formula  $F_i$  is quantified by a real value  $w_i$ , its weight. Together with a set of constants  $C$  it defines a Markov network  $M_{L,C}$  as follows:

- $M_{L,C}$  contains a node for each possible grounding of each predicate appearing in  $L$ .
- $M_{L,C}$  contains an edge between two nodes (i. e. ground atoms) iff the ground atoms appear together in at least one grounding of one formula in  $L$ .
- $M_{L,C}$  contains one feature (function) for each possible grounding of each formula  $F_i$  in  $L$ . The value of the feature for a possible world  $x$  is 1, if the ground formula is true for  $x$  (and 0 otherwise). Each feature is weighted by the weight  $w_i$  of its respecting formula  $F_i$ .

According to the above definition, a MLN (i. e. the weighted formulas) defines a template for constructing *ground Markov networks*. For a different set  $C'$  of constants, a different ground Markov network  $M_{L,C'}$  emerges from  $L$ . These ground Markov networks may vary in size, but their general structure is quite similar, e. g. the groundings of a formula  $F_i$  have the weight  $w_i$  in any ground Markov network of  $L$ . The ground Markov network  $M_{L,C}$  specifies

$$P(X = x) = \frac{1}{Z} \exp \left( \sum_i w_i n_i(x) \right)$$

as the probability distribution over possible worlds  $x$  (whereas  $Z$  is a normalization factor). For each formula  $F_i$ , the binary result values of its feature functions have been incorporated into the counting function  $n_i(x)$ , so  $n_i(x)$  compactly expresses the number of true groundings of  $F_i$  in the possible world  $x$ .

*Example 4.* In the following example, we model the relations described in Ex. 1 as a MLN (using the Alchemy syntax [15] for MLN files). The “!” operator

used in the predicate declarations of *lives\_in* and *neighborhood* enforces that the respective variables will have mutually exclusive and exhaustive values, i.e. that every person lives in exactly one town and one neighborhood (in terms of ground atoms). The weights of the formulas are to be understood exemplary (since “realistic” weights cannot be estimated just like that, but requires learning from data). We declare the typed predicates *alarm(person)*, *neighborhood(person, hood\_state!)*, *lives\_in(person, town!)*, *burglary(person)*, the types and constants *person* = {*James, Carl*}, *town* = {*Freiburg, Yorkshire, Austin*}, *hood\_state* = {*Bad, Average, Good*}, and add the following weighted formulas:

$$\begin{array}{lll}
2.2 & \textit{burglary}(x) & \Rightarrow \textit{alarm}(x) \\
2.2 & \textit{lives\_in}(x, y) \wedge \textit{tornado}(y) & \Rightarrow \textit{alarm}(x) \\
-0.8 & \textit{neighborhood}(x, \textit{Good}) & \Rightarrow \textit{burglary}(x) \\
-0.4 & \textit{neighborhood}(x, \textit{Average}) & \Rightarrow \textit{burglary}(x) \\
0.4 & \textit{neighborhood}(x, \textit{Bad}) & \Rightarrow \textit{burglary}(x)
\end{array}$$

### 2.3 Relational Maximum Entropy

In this paper we also consider a specific approach for reasoning under maximum entropy on first-order probabilistic conditional logic [11]. Knowledge is captured in RME using probabilistic conditionals as in probabilistic conditional logic, cf. [16].

**Definition 4 (RME conditional).** A RME conditional  $r = (\phi \mid \psi)[\alpha][c]$  consists of a head literal  $\phi$ , a list of  $n$  body literals  $\psi = \psi_1, \dots, \psi_n$ , a real value  $\alpha \in [0, 1]$ , and a list of meta-constraints  $c = c_1, \dots, c_m$ , which allows the restriction of the substitution for certain variables. A meta-constraint is either an expression of the form  $X \neq Y$  or  $X \notin \{k_1, \dots, k_l\}$ , with variables  $X, Y$  and  $\{k_1, \dots, k_l\} \subseteq U$ . A conditional  $r$  is called ground iff  $r$  contains no variables. The set of all RME conditionals is determined as the language  $(\mathcal{L} \mid \mathcal{L})^{\text{rel}}$  and the set of all ground conditionals is referred to by  $(\mathcal{L} \mid \mathcal{L})_U^{\text{rel}}$ .

**Definition 5 (RME knowledge base).** A RME knowledge base  $KB$  is a quadruple  $KB = (S, U, P, R)$  with a finite set of sorts  $S$ , a finite set of constants  $U$ , a finite set of predicates  $P$ , and a finite set of RME conditionals  $R$ . Any constant of the universe  $U$  is associated with one sort in  $S$  and  $U_\sigma$  determines the constants with sort  $\sigma$ . Each argument of a predicate  $p(\sigma_1, \dots, \sigma_k) \in P$  is also associated with a sort in  $S$ ,  $\sigma_i \in S, 1 \leq i \leq k$ . Furthermore, all variables  $\text{Var}(R)$  occurring in  $R$  are associated with a sort in  $S$ , too. Constants and variables are referred to as terms  $t$ .

*Example 5.* We represent Ex. 1 as a RME knowledge base  $KB$  which consists of sorts  $S = \{\textit{Person}, \textit{Town}, \textit{Status}\}$ , constants  $U_{\textit{Person}} = \{\textit{carl}, \textit{stefan}\}$  of sort *Person*,  $U_{\textit{Town}} = \{\textit{freiburg}, \textit{yorkshire}, \textit{austin}\}$  of sort *Town*,  $U_{\textit{Status}} = \{\textit{bad}, \textit{average}, \textit{good}\}$  of sort *Status*, predicates  $P = \{\textit{alarm}(\textit{Person}), \textit{burglary}(\textit{Person}), \textit{lives\_in}(\textit{Person}, \textit{Town}), \textit{neighbourhood}(\textit{Person}, \textit{Status})\}$ , and conditionals  $R = \{c_1, \dots, c_7\}$ .

$$\begin{aligned}
c_1 &= (\text{alarm}(\mathbf{X}) \mid \text{burglary}(\mathbf{X})) [0.9] \\
c_2 &= (\text{alarm}(\mathbf{X}) \mid \text{lives\_in}(\mathbf{X}, \mathbf{Y}), \text{tornado}(\mathbf{Y})) [0.9] \} \\
c_3 &= (\text{burglary}(\mathbf{X}) \mid \text{neighborhood}(\mathbf{X}, \text{bad})) [0.6] \\
c_4 &= (\text{burglary}(\mathbf{X}) \mid \text{neighborhood}(\mathbf{X}, \text{average})) [0.4] \\
c_5 &= (\text{burglary}(\mathbf{X}) \mid \text{neighborhood}(\mathbf{X}, \text{good})) [0.3] \\
c_6 &= (\text{neighborhood}(\mathbf{X}, \mathbf{Z}) \mid \text{neighborhood}(\mathbf{X}, \mathbf{Y})) [0.0] [\mathbf{Y} \neq \mathbf{Z}] \\
c_7 &= (\text{lives\_in}(\mathbf{X}, \mathbf{Z}) \mid \text{lives\_in}(\mathbf{X}, \mathbf{Y})) [0.0] [\mathbf{Y} \neq \mathbf{Z}]
\end{aligned}$$

Notice, that conditionals  $c_6$  and  $c_7$  ensure mutual exclusion of the states for literals of *neighborhood* and *lives\_in*.

Semantics are given to RME knowledge bases by grounding  $R$  with a *grounding operator* (GOP)<sup>6</sup> to a propositional probabilistic knowledge base and calculating the probability distribution with maximum entropy  $P_{\mathcal{G}_\chi(R)}^{ME}$ . A GOP (see Fig. 1) is a type of substitution pattern that facilitates the modeling of exceptional knowledge. For example, we could add the exceptional rule

$$c_8 = (\text{alarm}(\mathbf{X}) \mid \text{lives\_in}(\mathbf{X}, \text{freiburg}), \text{tornado}(\text{freiburg})) [0.1]$$

to  $KB$  in order to model that in Freiburg tornados are usually not strong enough to cause an alarm. Then the GOP  $\mathcal{G}_{\text{priority}}$  would prefer all instances of  $c_8$  above all instances of  $c_2$ , in which the constant *freiburg* occurs (for details see [11]).


After  $R$  is grounded,  $\mathcal{G}_\chi(R)$  is treated as a propositional knowledge base and  $P_{\mathcal{G}_\chi(R)}^{ME}$  can be calculated as in the propositional case [16]. A RME conditional  $Q \in (\mathcal{L} \mid \mathcal{L})^{rel}$  is *fulfilled under the grounding*  $\mathcal{G}_\chi(R)$  by the RME knowledge base  $R$  if the following hold:

$$R \models_{\mathcal{G}_\chi}^{ME} Q \iff P_{\mathcal{G}_\chi(R)}^{ME} \models \mathcal{G}_\chi(Q) \iff \forall q \in \mathcal{G}_\chi(Q) : P_{\mathcal{G}_\chi(R)}^{ME} \models q.$$

The RME inference process can be divided into three steps: 1.) ground the KB with a certain GOP  $\mathcal{G}_\chi$ , 2.) calculate the probability distribution  $P_{\mathcal{G}_\chi(R)}^{ME}$  with maximum entropy for the grounded instance  $\mathcal{G}_\chi(R)$ , and 3.) calculate all probabilistic implications of  $P_{\mathcal{G}_\chi(R)}^{ME}$ .

A more elaborated overview on the framework of relational maximum entropy as introduced here is given in [11].

<sup>6</sup> In [11] the *naive*-, *cautious*-, *conservative*-, and *priority*-grounding strategies are presented and analyzed.



**Fig. 1.** RME semantics overview. A RME KB  $R$  is transformed into a propositional representation by the GOP  $\mathcal{G}_\chi$ . The result  $\mathcal{G}_\chi(R)$  is used to calculate the Max-Ent-distribution  $P_{\mathcal{G}_\chi(R)}^{ME}$  in order to answer the ground query  $\mathcal{G}_\chi(q)$ .

### 3 System Architecture

KREATOR<sup>7</sup> is an integrated development environment for representing, reasoning, and learning with relational probabilistic knowledge. Still being in development KREATOR aims to become a versatile toolbox for researchers and knowledge engineers in the field of statistical relational learning. KREATOR is written in Java and thus is designed using the object-oriented programming paradigm. It facilitates several architectural and design patterns such as model-view control, abstract factories, and command patterns. Central aspects of the design of KREATOR are *modularity*, *extensibility*, *usability*, *reproducibility*, and its intended application in scientific research.

*Modularity and Extensibility* KREATOR is modular and extensible with respect to several components. In the following we discuss just two important aspects. First, KREATOR separates between the internal logic and the user interface using an abstract command structure. Each top-level functionality of KREATOR is internally represented and encapsulated in an abstract `KReatorCommand`. Consequently, the user interface can be exchanged or modified in an easy and unproblematic way, because it is separated from the internal program structure by this `KReatorCommand` layer. As a matter of fact, the current version of KREATOR features both a graphical user interface and a command line interface (the KREATOR *console*) which processes commands in KREATORSCRIPT syntax (see Sec. 4.1). Second, KREATOR was designed to support many different approaches for relational knowledge representation, cf. Sec. 2. As a consequence, KREATOR features very abstract notions of concepts like knowledge bases, queries and data sets that can be implemented by a specific approach. At the moment, KREATOR supports knowledge representation using Bayesian logic programs (BLPs), Markov logic networks (MLNs), and the relational maximum entropy (RME) approach described in Sec. 2.3. Other formalisms will be integrated in the near future.

<sup>7</sup> The “KR” in KREATOR stands for “Knowledge Representation” and the name KREATOR indicates its intended usage as a development environment for knowledge engineers.



*Usability and Reproducibility* An important design aspect of KREATOR and especially of the graphical user interface is usability. While prototypical implementations of specific approaches to relational probabilistic knowledge representation (and approaches for any problem in general) are essential for validating results and evaluation, these software solutions are often very hard to use and differ significantly in their usage. Especially when one wants to compare different solutions these tools do not offer an easy access for new users. KREATOR features a common and simple interface to different approaches of relational probabilistic knowledge representation within a single application.

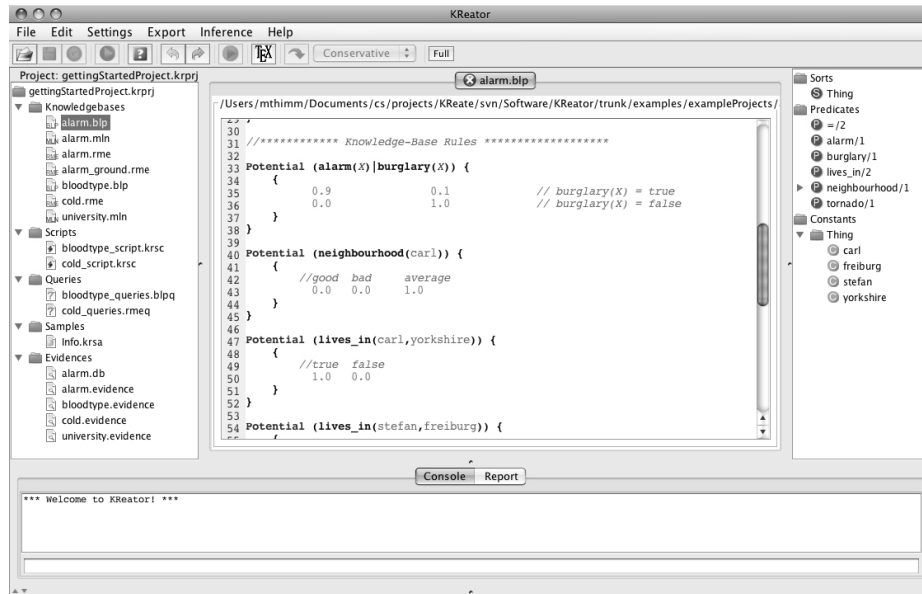
*Application in Scientific Research* Both usability and reproducibility are important aspects when designing a tool for conducting scientific research. Besides that, other important features are also provided within KREATOR. For example, KREATOR can export knowledge base files as formatted L<sup>A</sup>T<sub>E</sub>X output, making the seamless processing of example knowledge bases in scientific publications very convenient. KREATOR records every user operation (no matter whether it was caused by GUI interaction or by a console input) and its result in a *report*. Since all operations are reported in KREATORSCRIPT syntax, the report itself represents a valid script. Therefore the whole report or parts of it can be saved as a KREATORSCRIPT file which can be executed anytime to repeat the recorded operations.

*Used frameworks* KREATOR makes use of well-established software frameworks to process some of the supported knowledge representation formalisms. Performing inference on MLNs is handled entirely by the Alchemy software package [15], a console-based tool for processing Markov logic networks. Alchemy is open source software developed by the inventors of Markov logic networks and can freely be obtained on <http://alchemy.cs.washington.edu/>. To process ground RME knowledge bases, an appropriate reasoner for maximum entropy must be utilized. KREATOR does not directly interact with a certain reasoner. Instead, KREATOR uses a so-called ME-adapter to communicate with a (quite arbitrary) MaxEnt-reasoner. Currently, such an adapter is supplied for the SPIRIT reasoner [16]. SPIRIT is a tool for processing (propositional) conditional probabilistic knowledge bases using maximum entropy methods and can be obtained on [http://www.fernuni-hagen.de/BWLOR/spirit\\_int/](http://www.fernuni-hagen.de/BWLOR/spirit_int/). An appropriate adapter for the MECORE reasoner [17] has also been developed.

## 4 Usage

KREATOR comes with a graphical user interface and an integrated console-based interface. The main view of KREATOR (see Fig. 2) is divided into the menu and toolbars and four main panels: the project panel, the editor panel, the outline panel, and the console panel.

*The project panel* KREATOR structures its data into projects which may contain knowledge bases, scripts written in KREATORSCRIPT (see below), query collections for knowledge bases, and sample/evidence files. Although all types of files



**Fig. 2.** KReator – Main window

can be opened independently in KREATOR, projects can help the knowledge engineer to organize his work. The project panel of KREATOR (seen in the upper left in Fig. 2) gives a complete overview on the project the user is currently working on.

*The editor panel* All files supported by KREATOR can be viewed and edited in the editor panel (seen in the upper middle in Fig. 2). Multiple files can be opened at the same time and the editor supports editing knowledge bases and the like with syntax-highlighting, syntax check, and other features normally known from development environments for programming languages.

*The outline panel* The outline panel (seen in the upper right in Fig. 2) gives an overview on the currently viewed file in the editor panel. If the file is a knowledge base the outline shows information on the logical components of the knowledge base, such as used predicates (and, in case of BLPs, their states), constants, and sorts (if the knowledge base uses a typed language).

*The console panel* The console panel (seen at the bottom in Fig. 2) contains two tabs, one with the actual console interface and one with the *report*. The console can be used to access nearly every KREATOR functionality just using textual commands, e. g. querying knowledge bases, open and saving file, and so on. The console is a live interpreter for KREATORSCRIPT, the scripting language also used for writing scripts (see below). The console panel also contains the report tab. Every action executed in KREATOR, e. g. opening a file in the graphical

user interface or querying a knowledge base from the console, is recorded as a KREATORSCRIPT command in the report. The whole report or parts of it can easily be saved as script file and executed again when experiments have to be repeated and results have to be reproduced.

#### 4.1 The KReatorScript Language

The KREATORSCRIPT language incorporates all those commands which can be processed by the console and by script files as well (see Fig. 3 for some KREATORSCRIPT lines). Since there are commands available for all high-level KREATOR functionalities, every sequence of working steps can be expressed as an appropriate command sequence in a KREATORSCRIPT file. Thus, the (re-)utilization of scripts can clearly increase the efficiency and productivity when working with KREATOR. As mentioned above, the input console and report work hand in hand with KREATOR's scripting functionality, making the KREATORSCRIPT language a strong instrument in the whole working process.

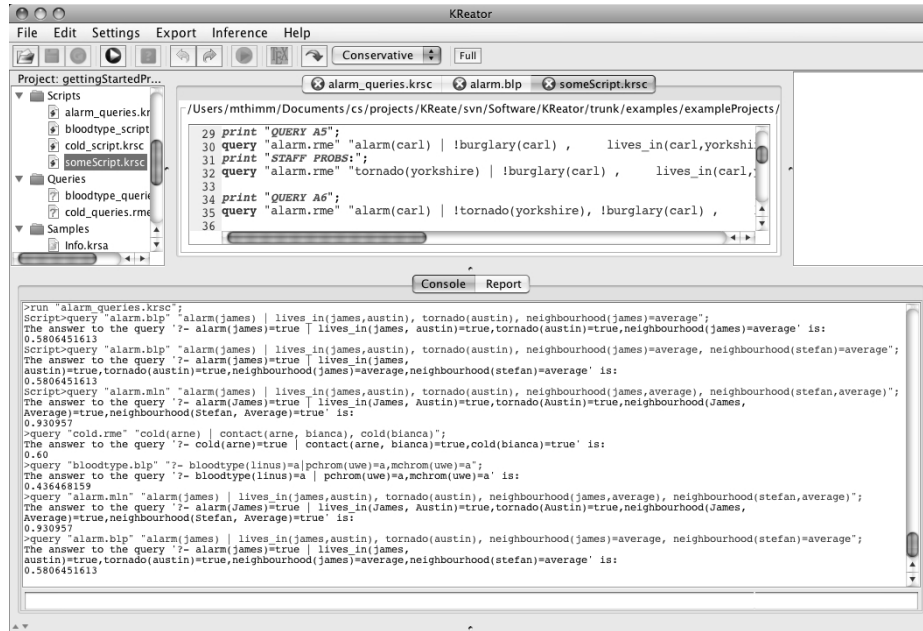
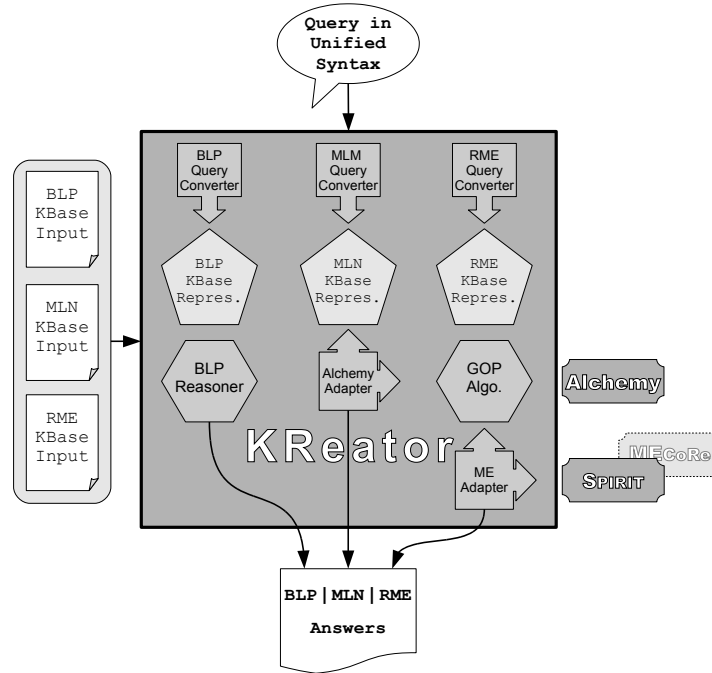


Fig. 3. The console with some KREATORSCRIPT

#### 4.2 Querying a Knowledge Base

One of the most important tasks when working with knowledge bases is to address queries to a knowledge base, i.e. to infer knowledge. For that reason, KREATOR provides several functionalities which simplify the dealing with queries and make it more efficient.

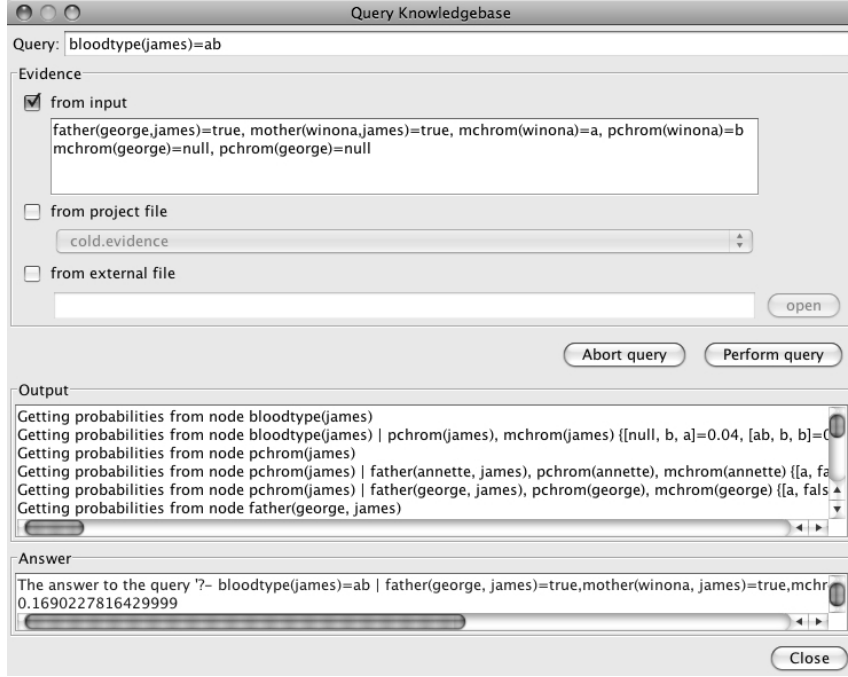
KREATOR permits the processing of queries expressed in a *unified query syntax*. This query syntax abstracts from the respective syntax which is necessary to address a “native” query to a BLP, MLN, or RME knowledge base (and which also depends on the respective inference engine). That way, a query in unified syntax can be passed to an appropriate BLP, MLN, and RME knowledge base as well. The idea behind this functionality is, that some knowledge (cf. Ex. 1) can be modeled in different knowledge representation approaches (cf. Ex. 2, Ex. 4, and Ex. 5) and the user is able to compare these approaches in a more direct way. Such a comparison can then be done by formulating appropriate queries in unified syntax, passing them to the different knowledge bases, and finally analyzing the different answers, i.e. the probabilities. A KREATOR query in unified syntax consists of two parts: In the “head” of the query there are one or more ground atoms whose probabilities shall be determined. The “body” of the query is composed of several evidence atoms. For each supported knowledge representation formalism, KREATOR must convert a query in unified syntax in the exact syntax required by the respective inference engine. KREATOR also converts the respective output results to present them in a standardized format to the user. Figure 4 illustrates the processing of a query in unified syntax.



**Fig. 4.** Processing query in unified syntax

KREATOR offers the user an easy way to address a query to a knowledge base, simply by calling its query dialogue. In this dialogue (Fig. 5), the user can

input the atoms to be queried and he can conveniently specify the evidence. Since evidence usually consists of several atoms and is often reused for different queries, the user has the option to specify a file which contains the evidence to be considered. While processing a query, the output area of the dialogue informs about important steps of the inference process. The calculated answer is clearly displayed in the lower part of the dialogue.



**Fig. 5.** Querying a knowledge base

Besides the capability of passing individual (i. e. ad-hoc) queries to a knowledge base, KREATOR also supports so-called *query collections*. A query collections is a file which contains several queries (either in unified or native syntax). Such a query collection can be passed to a knowledge base, so that all included queries are processed one after another. That way, KREATOR supports a persistent handling and batch-style processing of queries.

*Example 6.* Continuing our previously introduced burglary example consider the following evidence:

*lives\_in(james, yorkshire), lives\_in(stefan, freiburg), burglary(james),  
tornado(freiburg), neighborhood(james) = average, neighborhood(stefan) = bad*

Table 1 shows three queries and their respective probabilities inferred from each of the example knowledge bases. The nine separate calculations altogether took

about three seconds. Each of the three knowledge bases represents Ex. 1 by a different knowledge representation approach. Nevertheless, the inferred probabilities are quite similar, except for the significant lower BLP probability of the query *alarm(stefan)*. This results from using noisy-or as combining rule in the BLP calculations and from the fact that the CPDs of the BLP knowledge base carry some information not incorporated in the MLN and BLP knowledge bases.

	BLP	MLN	RME
<i>alarm(james)</i>	0.900	0.896	0.918
<i>alarm(stefan)</i>	0.650	0.896	0.907
<i>burglary(stefan)</i>	0.300	0.254	0.354

**Table 1.** Exemplary queries on different representations of the burglary example.

## 5 Summary and Future Work

In this paper we presented KREATOR and illustrated its system architecture and usage. Although KREATOR is still in an early stage of development it already supports Bayesian logic programs, Markov logic networks, and relational maximum entropy. Thus KREATOR is a versatile toolbox for probabilistic relational reasoning and alleviates the researcher’s and knowledge engineer’s work with different approaches to statistical relational learning.

Since KREATOR is still in an early development stage, there are a lot of plans on future development. The integration of adequate learning algorithms will be one of our major tasks in the near future, as our main focus so far was the integration of reasoning components. We also plan to integrate an extended version of the CONDOR system [18] and other formalisms for relational probabilistic knowledge representation such as logical Bayesian networks [7] and probabilistic relational models [9], as well as to use KREATOR as a testbed to evaluate other approaches for relational probabilistic reasoning under maximum entropy.

We plan to enhance KREATOR’s unified query syntax to allow more complex queries. This requires more sophisticated conversion patterns to translate a unified query to the respective target syntax, e.g. to handle multi-state BLP predicates in an automated way. The enhancement of the query syntax will go along with the development of an even more challenging feature: We plan on introducing some kind of unified knowledge base (template) format. The final goal is to be able to formulate (at least) the central aspects of a knowledge base in a unified syntax and to have this knowledge base be converted to different target languages (at least semi-)automatically. Having this functionality available would dramatically improve the handling and comparison of different knowledge representation formalisms.

KREATOR is available under the GNU General Public License and can be obtained from <http://ls6-www.cs.uni-dortmund.de/kreator/>. The future development of KREATOR will also be strongly influenced by the feedback of early users. We will include such feedback in our development decisions and try to prioritize such aspects which are most important to the users.

**Acknowledgements.** The research reported here was partially supported by the Deutsche Forschungsgemeinschaft (grants BE 1700/7-1 and KE 1413/2-1).

## References

1. De Raedt, L., Kersting, K.: Probabilistic Inductive Logic Programming. In De Raedt, L., Kersting, K., Landwehr, N., Muggleton, S., Chen, J., eds.: Probabilistic Inductive Logic Programming. Springer (2008) 1–27
2. Cussens, J.: Logic-based Formalisms for Statistical Relational Learning. In Getoor, L., Taskar, B., eds.: An Introduction to Statistical Relational Learning. MIT Press (2007)
3. Kersting, K., De Raedt, L.: Bayesian Logic Programming: Theory and Tool. In Getoor, L., Taskar, B., eds.: An Introduction to Statistical Relational Learning. MIT Press (2007)
4. Domingos, P., Richardson, M.: Markov Logic: A Unifying Framework for Statistical Relational Learning. In: Proceedings of the ICML-2004 Workshop on Statistical Relational Learning and its Connections to Other Fields. (2004) 49–54
5. Pearl, J.: Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. Morgan Kaufmann (1998)
6. Wellman, M.P., Breese, J.S., Goldman, R.P.: From Knowledge Bases to Decision Models. The Knowledge Engineering Review **7**(1) (1992) 35–53
7. Fierens, D.: Learning Directed Probabilistic Logical Models from Relational Data. PhD thesis, Katholieke Universiteit Leuven (2008)
8. Jaeger, M.: Relational Bayesian Networks: A Survey. Electronic Transactions in Artificial Intelligence **6** (2002)
9. Getoor, L., Friedman, N., Koller, D., Taskar, B.: Learning Probabilistic Models of Relational Structure. In Brodley, C.E., Danyluk, A.P., eds.: Proc. of the 18th International Conf. on Machine Learning (ICML 2001), Morgan Kaufmann (2001)
10. Baral, C., Gelfond, M., Rushton, N.: Probabilistic Reasoning with Answer Sets. Theory and Practice of Logic Programming (2009)
11. Loh, S.: Relational Probabilistic Inference Based on Maximum Entropy. Master’s thesis, Technische Universität Dortmund (2009)
12. Thimm, M.: Representing Statistical Information and Degrees of Belief in First-Order Probabilistic Conditional Logic. In: Proceedings of the First Workshop on Relational Approaches to Knowledge Representation and Learning, Paderborn, Germany (September 2009)
13. Ketkar, N.S., Holder, L.B., Cook, D.J.: Comparison of graph-based and logic-based multi-relational data mining. SIGKDD Explor. Newsl. **7**(2) (2005) 64–71
14. Muggleton, S., Chen, J.: A Behavioral Comparison of some Probabilistic Logic Models. In Raedt, L.D., Kersting, K., Landwehr, N., Muggleton, S., Chen, J., eds.: Probabilistic Inductive Logic Programming. Springer (2008) 305–324
15. Kok, S., Singla, P., Richardson, M., Domingos, P., Sumner, M., Poon, H., Lowd, D., Wang, J.: The Alchemy System for Statistical Relational AI: User Manual. Department of Computer Science and Engineering, University of Washington. (2008)
16. Rödder, W., Meyer, C.H.: Coherent Knowledge Processing at Maximum Entropy by SPIRIT. In: Proceedings UAI 1996. (1996) 470–476
17. Finthammer, M., Beierle, C., Berger, B., Kern-Isberner, G.: Probabilistic Reasoning at Optimum Entropy with the MEcore System. In: Proceedings of FLAIRS’09, AAAI Press (2009)
18. Fisseler, J., Kern-Isberner, G., Beierle, C., Koch, A., Moeller, C.: Algebraic Knowledge Discovery using Haskell. In: Practical Aspects of Declarative Languages, 9th International Symposium. Springer (2007)