

MV-Datalog+-: Effective Rule-based Reasoning with Uncertain Observations

MATTHIAS LANZINGER

University of Oxford

STEFANO SFERRAZZA

University of Oxford

Technische Universität Wien

GEORG GOTTLÖB

University of Oxford

Abstract

Modern applications combine information from a great variety of sources. Oftentimes, some of these sources, like Machine-Learning systems, are not strictly binary but associated with some degree of (lack of) confidence in the observation.

We propose MV-Datalog and MV-Datalog[±] as extensions of Datalog and Datalog[±], respectively, to the fuzzy semantics of infinite-valued Lukasiewicz logic L as languages for effectively reasoning in scenarios where such uncertain observations occur. We show that the semantics of MV-Datalog exhibits similar model theoretic properties as Datalog. In particular, we show that (fuzzy) entailment can be decided via minimal fuzzy models. We show that when they exist, such minimal fuzzy models are unique and can be characterised in terms of a linear optimisation problem over the output of a fixed-point procedure. On the basis of this characterisation, we propose similar many-valued semantics for rules with existential quantification in the head, extending Datalog[±]. This paper is under consideration for acceptance in TPLP.

KEYWORDS: Artificial intelligence, machine learning, Datalog, fuzzy logic, uncertainty, Datalog[±],

1 Introduction

Datalog and its extensions have long been important languages in databases and at the foundation of many rule-based reasoning formalisms. Continuous theoretical and technical improvements have led to the development of highly efficient Datalog systems for widespread practical use in a variety of applications (see e.g, Maier et al. (2018)).

However, in many such real-world applications the observations and facts that serve as the input are not actually certain but rather associated with some (possibly unknown) level of uncertainty. One particularly important situation where such a setting occurs are settings where observations are made by Machine Learning (ML) systems.

For example, say a database contains a relation that labels images of animals with the class of animal observed in the image. In modern settings, such labels are commonly derived from the output of ML systems that attempt to classify the image contents. Generally, such systems output possible labels for an image, together with a score that

can be interpreted as the level of confidence in the labels correctness (cf., Mishkin et al. (2017); Lee et al. (2017)).

When the labelling in our example is used in further logical inference, the level of confidence in the individual labels would ideally also be reflected in any knowledge inferred from these uncertain observations. However, classical logical reasoning provides no way to consider such information and requires the projection of uncertainty levels to be either true or false, usually via some simple numerical threshold. Beyond the overall loss of information, this can also lead to problematic outcomes where logical conclusions are true despite their derivation being based on observations made with only moderate confidence. With observations from ML systems becoming a commonplace feature of many modern reasoning settings, the projection to binary truth values is severely limiting the immense potential of integrating ML observations with logic programming.

Related Work There are two natural ways to interpret uncertain observations as described in the example above. One can consider them as the likelihood of the fact being true, i.e., the level of confidence in a fact is interpreted as a probability of the fact holding true. Such an interpretation has been widely studied in the context of Problog (Raedt et al. 2007), Markov Logic Networks (Richardson and Domingos 2006), and probabilistic databases (Suciu et al. 2011; Cavallo and Pittarelli 1987). However, generally such formalisms make strong assumptions on the pairwise probabilistic independence of all tuples which can be difficult to satisfy. An approach to probabilistic reasoning in Datalog[±] based on the chase procedure was proposed very recently by Bellomarini et al. (2020).

Alternatively, one can express levels of confidence in a fact in terms of *degrees of truth* as in *fuzzy logics* (cf. Hájek (1998)). That is, a fact that is considered to be true to a certain degree, in accordance with the level of confidence in the observation.

There have been a wide variety of approaches of combining logic programming with fuzzy logic. Pioneering work in this area by Achs and Kiss (1995), Ebrahim (2001), and Vojtás (2001) introduced a number of early foundational ideas for fuzzy logic programming. The first also considers fuzziness in the object domain (expressing the possibility that two constants may actually refer to the same object), via fuzzy similarity relations on the object domain. This idea was further expanded in the context of deductive databases by Iranzo and Sáenz-Pérez (2018).

One limiting factor of these approaches in our intended setting is the reliance on the *Gödel t-norm* (see, Preining (2010)) as the basis for many-valued semantics. The Gödel t-norm of two formulas ϕ , ψ is the minimum of the respective truth degrees of ϕ and ψ . While this simplifies the model theory of the resulting logics, the resulting semantics are not ideal for our envisioned applications in which we want to combine uncertain observations from AI systems with large certain knowledge bases. Further discussion of the important differences between semantics based on the Gödel t-norm and the Lukasiewicz t-norm (see, Hájek (1998)), which we propose to use in this paper, is provided in Section 2.

Finally, Probabilistic Soft Logic (PSL) was introduced by Bach et al. (2017) as a framework for fuzzy reasoning with semantics based on Łukasiewicz logic. PSL is a significantly different language with different use cases than our proposed languages. It allows for negation in rule bodies as well as disjunction in the head, in addition to certain forms of aggregation. Semantically it treats all rules as soft constraints and aims to

find interpretations that satisfy rules as much as possible. Similar ideas have also been proposed in fuzzy description logics, see e.g., Bobillo and Straccia (2016); Borgwardt and Peñaloza (2017); Stoilos et al. (2007), or Lukasiewicz and Straccia (2008) for an overview of fuzzy DLs.

Contributions We propose extensions of Datalog and Datalog[±], i.e., Datalog with existential quantification in the head, to the many-valued semantics of infinite-valued Lukasiewicz logic. We argue that these languages offer natural semantics for practical applications with uncertain data while preserving key theoretical properties of the classical (two-valued) case. Overall, the contributions of this paper are summarised as follows.

- We introduce MV-Datalog as Datalog with Lukasiewicz semantics. We discuss adaptations of the standard minimal model concept to fuzzy models where every fact takes the least possible truth degree. We show that for consistent inputs, unique minimal models always exist and give a characterisation of such minimal models. Our results hold also for K -fuzzy models, i.e., interpretations that satisfy formulas of the program only to degree K rather than fully.
- We show that MV-Datalog is suitable for practical applications. For one, our characterisation reveals that finding K -fuzzy models is feasible by building on standard methods as from the classical setting. In particular, we also show fuzzy fact entailment in MV-Datalog has the same complexity as fact entailment in Datalog.
- We propose MV-Datalog[±] as practice oriented extension of MV-Datalog by existential quantification in the head. To obtain intuitive semantics for logic programming we deviate from standard semantics of many valued existential quantification and instead argue for the use of strong existential quantification. We discuss fundamental model theoretic properties and challenges of MV-Datalog[±] and show how our characterisation result for MV-Datalog can be extended to capture a natural class of preferred models of MV-Datalog[±], providing the foundation for a system implementation of MV-Datalog[±].

Structure The rest of this paper is structured as follows. Terminology, notation and further preliminaries are stated in Section 2. Throughout Section 3, we introduce MV-Datalog and our main results for minimal fuzzy models in MV-Datalog. The further extension to MV-Datalog[±] is presented in Section 4. We conclude with final remarks and a brief discussion of future work in Section 5. Proof details that had to be omitted from the main text due to space restrictions are provided in Appendix A.

2 Preliminaries

We assume familiarity with standard notions of Datalog and first-order logic. Throughout this document we consider only logical languages with no function symbols.

Datalog Fix a relational signature σ and a (countable) domain Dom . A (*relational*) *atom* is a term $R(u_1, \dots, u_{\#R})$ where $\#R$ is the arity of relation symbol $R \in \sigma$ and for $0 \leq i \leq \#R$, u_i is a variable or a constant. If all positions of the atom are constants we say that it is a *ground atom* (or *fact*). A (*Datalog*) *rule* r is a first-order implication $\varphi \rightarrow \psi$ where

φ is a conjunction of relational atoms and ψ consists of a single relational atom. We say that φ is the *body* of r ($\text{body}(r)$) and ψ is the *head* of r ($\text{head}(r)$). A rule is *safe* if all free variables in the head of a rule also occur in its body. A *Datalog program* is a set of safe Datalog rules. The models of a Datalog program are the first-order interpretations that satisfy all rules. For two models $\mathcal{I}, \mathcal{I}'$ of a Datalog program, we say that $\mathcal{I} \leq \mathcal{I}'$ if every every ground atom satisfied by \mathcal{I} is also satisfied by \mathcal{I}' . The *minimal model* of a Datalog program Π is the model \mathcal{I} of Π such that for all other models \mathcal{I}' of Π it holds that $\mathcal{I} \leq \mathcal{I}'$.

A *database* D is a set of ground atoms, logically interpreted as a theory consisting of an atomic sentence for each ground atom. We say that \mathcal{I} is a model of a program and a database (Π, D) , if \mathcal{I} is a model of both Π and D . We write $\Pi, D \models \varphi$ if formula φ is satisfied for every model of (Π, D) . Note that this is equivalent to satisfaction under the minimal model of (Π, D) .

Datalog[±] The extension of Datalog with existential quantification in the head, led to the well studied Datalog[±] Cali et al. (2012). A *Datalog[±] rule* is either a Datalog rule or a first-order formula of the form $\varphi \rightarrow \exists \mathbf{y} \psi(\mathbf{y})$ where φ is again a conjunction of relational atoms and ψ is a relational atom that mentions all variables in \mathbf{y} . A *Datalog[±] program* is a set of safe Datalog[±] rules. Models in Datalog[±] are defined in the same way as for plain Datalog. Note that existential quantification ranges over the whole (infinite) domain Dom rather than the active domain of the program/database. It follows that, in contrast to Datalog, Datalog[±] programs do not have a unique minimal model. Furthermore, fact entailment in Datalog[±] is undecidable in general but vast decidable fragments are known.

Oblivious Chase Fact entailment in Datalog[±] (and Datalog) is closely linked to the so-called *chase* fixed-point procedures. In the following we will be particularly interested in a specific chase procedure called the *oblivious chase*. We give only the minimal necessary technical details here and refer to Cali et al. (2013) for full details.

For any formula φ , we write $\text{vars}(\varphi)$ for the set of *free variables* in the formula. Given a database D , we say that a Datalog[±] rule r is *obliviously applicable* to D if there exists a homomorphism $h: \text{vars}(r) \rightarrow \text{Dom}$ such that all atoms in $h(\text{body}(r))$ are in D .

Let r be a Datalog[±] rule that is obliviously applicable to D via a homomorphism h . Let h' be the extension of h to the existentially quantified variables \mathbf{y} in $\text{head}(r)$ such that for each $y \in \mathbf{y}$, $h'(y)$ is a fresh constant, called a *labelled null*, not occurring in D ($h' = h$ when r is a Datalog rule). Let ψ be the atomic formula in $\text{head}(r)$. The result of the *oblivious application* of r, h to D is $D' = D \cup h'(\psi)$. We write $D \xrightarrow{r, h} D'$. Let D be a database and Π a Datalog[±] program. The *oblivious chase sequence* $\text{OChase}(\Pi, D)$ is the sequence D_0, D_1, D_2, \dots with $D = D_0$ such that for all $i \geq 0$, $D_i \xrightarrow{r_i, h_i} D_{i+1}$ and $r_i \in \Pi$, the chase sequence proceeds as long as there are oblivious applications r, h that have not yet been applied. We refer to the limit of the oblivious chase sequence as $\text{OLim}(\Pi, D)$. Every oblivious application r, h induces a grounding of r where all variables (including the existentially quantified ones) are assigned constants according to h and any existential quantification is removed. We refer to the set of all ground rules induced in this way by the applications in $\text{OChase}(\Pi, D)$ as $\text{OGround}(\Pi, D)$. Note that for every Π and D , $\text{OLim}(\Pi, D)$ is unique up to isomorphism.

Lukasiewicz Logic While there are many semantics for many-valued logic, in this paper we will consider infinite-valued Lukasiewicz logic L (see, Hájek (1998)). Some of our techniques are particular to Lukasiewicz logic and the results given in this paper generally do not apply to other many-valued logics. For some relational signature σ , we consider the following logical language, where R is a relational atom (in σ) and a *formula* φ is defined via the grammar

$$\varphi ::= R \mid \varphi \odot \varphi \mid \varphi \oplus \varphi \mid \varphi \rightarrow \varphi \mid \neg \varphi$$

For a signature σ and domain Dom , let $GAtoms$ be the set of all ground atoms with respect to σ and Dom . A *truth assignment* is a function $\nu: GAtoms \rightarrow [0, 1]$, intuitively assigning a degree of truth in the real interval $[0, 1]$ to every ground atom. By slight abuse of notation we also apply ν to formulas to express the truth of ground formulas γ , γ' according to the following inductive definitions.

$$\begin{aligned} \nu(\neg \gamma) &= 1 - \nu(\gamma) \\ \nu(\gamma \odot \gamma') &= \max\{0, \nu(\gamma) + \nu(\gamma') - 1\} \\ \nu(\gamma \oplus \gamma') &= \min\{1, \nu(\gamma) + \nu(\gamma')\} \\ \nu(\gamma \rightarrow \gamma') &= \min\{1, 1 - \nu(\gamma) + \nu(\gamma')\} \end{aligned}$$

That is, \odot is the usual Lukasiewicz *t-norm* and \oplus is the corresponding *t-conorm*, which take the place of conjunction and disjunction, respectively. Implication $\gamma \rightarrow \gamma'$ could equivalently be defined as usual as $\neg \gamma \oplus \gamma'$. Note that De Morgan's rule holds as usual for \odot and \oplus .

For rational $K \in (0, 1]$ we say that a formula φ is *K-satisfied* by ν if for every grounding γ of φ over Dom it holds that $\nu(\gamma) \geq K$. Whenever we make use of K in this context we assume it to be rational. In the context of rule-based reasoning it may be of particular interest to observe that an implication is 1-satisfied exactly when the head is at least as true as the body. For a set of formulas Π , we say that a truth assignment ν is a *K-fuzzy model* if all formulas in Π are K -satisfied by ν .

In place of the database in the classical setting, we instead consider (*finite*) *partial truth assignments*, that is, partial functions $\tau: GAtoms \rightarrow (0, 1]$ that are defined for a finite number of ground atoms. Let (Π, τ) be a pair where Π is a set of formulas and τ is a partial truth assignment, a *K-fuzzy model* of (Π, τ) is a K -fuzzy model ν of Π where $\nu(G) = \tau(G)$ for every ground atom G for which τ is defined. Whenever we talk about formulas Π and partial truth assignments τ we use $ADom$ to refer to their *active domain*, i.e., the subset of the domain that is mentioned in either Π or $\{G \in GAtoms \mid \tau(G) \text{ is defined}\}$. We write $GAtoms[Adom]$ to indicate $GAtoms$ restricted to groundings over $Adom$.

K-fuzzy models and their theory have been introduced and studied in the context of logic programming by Ebrahim (2001) but under different many-valued semantics based on the Gödel t-norm.

Example 1. Consider a labelling of images as described in the introduction, represented by predicate *Label*. Additionally, we have a predicate corresponding to whether the image was taken in a polar region. Suppose we have the following observations for image i_1 .

$$\tau(\text{Label}(i_1, \text{Whale})) = 0.8 \quad \tau(\text{PolarRegion}(i_1)) = 0.7$$

Now we consider satisfaction of the following ground rule under the semantics studied by Ebrahim (2001) in contrast to Lukasiewicz semantics.

$$\text{Label}(i_1, \text{Whale}) \odot \text{PolarRegion}(i_1) \rightarrow \text{Orca}(i_1)$$

If we interpret implication and \odot under the semantics proposed by Ebrahim (2001), the body has truth 0.7 (the Gödel t-norm of 0.8 and 0.7). The implication evaluates to the maximum of 0.3 ($= 1 - \text{truth of the body}$) and the truth of the head. Hence, a model that 1-satisfies the rule would have to satisfy $\text{Orca}(i_1)$ with truth degree 1. This can be unintuitive, especially in the context of logic programming, since the truth degree of the fact is not inferred from the truth of the body.

In Lukasiewicz semantics, the body has truth degree 0.5 following the intuition that the conjunction of two uncertain observations becomes even less certain. For the implication to be 1-satisfied it suffices if $\text{Orca}(i_1)$ is true to at least degree 0.5. We see that the inferred truth degree of the head naturally depends on the truth of the body.

3 MV-Datalog and Minimal Fuzzy Models

In this section we propose the extension of Datalog to Lukasiewicz semantics for many-valued reasoning. We show that MV-Datalog allows for an analogue of minimal model semantics in the fuzzy setting. Moreover, we give a characterisation of such minimal models in terms of an optimisation problem over the ground program corresponding to the oblivious chase.

3.1 MV-Datalog

Definition 1 (MV-Datalog Program). *An MV-Datalog program Π is a set of L formulas of the form*

$$B_1 \odot B_2 \odot \dots \odot B_n \rightarrow H$$

where all B_i , for $1 \leq i \leq n$, and H are relational atoms.

As noted in Section 2, we consider a partial truth assignment τ in the place of a database. We will therefore also refer to τ as the database. We call a pair Π, τ of an MV-Datalog program and database an *MV-instance*. We can map a MV-Datalog program Π naturally to a Datalog program by substituting \odot by \wedge . We refer to the resulting Datalog program as Π^{crisp} . For the respective crisp version of τ we write D_τ for the database containing all facts G for which $\tau(G)$ is defined and greater than 0.

Analogous to fact entailment in classical Datalog, we have the central decision problem K-TRUTH, whether a fact is true to at least a degree c in all models, as follows.

K-TRUTH

Input MV-instance (Π, τ) , ground atom G , $c \in [0, 1]$

Output $\nu(G) \geq c$ for all K -fuzzy models ν of (Π, τ) ?

MV-Datalog is a proper extension of Datalog in the sense that for $K = 1$ and when all ground atoms in the database are assigned truth 1, its models coincide exactly with those of Datalog programs.

Note that in contrast to Datalog, MV-instances do not always have models. For example, consider a program consisting only of rule $R(x) \rightarrow S(x)$. For any $K > 0$, there is a database τ that assigns truth 1 to $R(a)$ and some truth less than K to $S(a)$ such that the rule is not K -satisfied. This is a consequence of the definition of a K -fuzzy model ν of (Π, τ) , which requires the truth values in ν to agree exactly with τ , for every fact for which τ is defined. In some settings it may be desirable to relax this slightly and consider K -fuzzy models ν where $\nu(G) \geq \tau(G)$ where $\tau(G)$ is defined¹. Our semantics cover such a relaxation since it can be simulated by straightforward rewriting of the program: for every relation symbol R that occurs in τ add rule $R(\mathbf{x}) \rightarrow R'(\mathbf{x})$, and replace all other occurrences of R in the program by R' .

Ebrahim (2001) introduced minimal K -fuzzy models to study logic programming with the Gödel t-norm, as the intersection of all K -fuzzy models, where truth values of an intersection are taken to the minimum between the two values. We will use the following alternative (but equivalent) definition here that will be more convenient in our setting. For two truth assignment ν, μ we write $\nu \leq \mu$ when for all $G \in GAtoms$, it holds that $\nu(G) \leq \mu(G)$. We similarly write $\nu < \mu$ if $\nu \leq \mu$ and $\nu(G) < \mu(G)$ for at least one $G \in GAtoms$.

Definition 2. Let Π, τ be an MV-instance. We say that a K -fuzzy model μ of (Π, τ) is minimal if for every K -fuzzy model ν of (Π, τ) it holds that $\mu \leq \nu$.

Minimal K -fuzzy models behave very similar to minimal models in Datalog. In particular, the two following propositions state two key properties allows us to focus on a single minimal fuzzy model for the problem K-TRUTH.

Theorem 1. Let Π, τ be an MV-instance. For every rational $K \in (0, 1]$, if (Π, τ) is K -satisfiable, then there exists a unique minimal K -fuzzy model for (Π, τ) .

Proof. We show that for every set of K -fuzzy models S , the infimum of S , defined as the truth assignment ν' with $\nu'(G) = \inf_{\nu \in S} \nu(G)$ for all $G \in GAtoms$, is also a K -fuzzy model for (Π, τ) .

For any ground rule γ and $\nu \in S$, the body truth value $\nu(\text{body}(\gamma))$ is given by $\max\{0, \left(\sum_{i=1}^{\ell} \nu(G_i)\right) - \ell + 1\}$ where the body of γ consists of ground atoms G_1, \dots, G_{ℓ} . Hence, since we have $\nu'(G_i) \leq \nu(G_i)$ for every $1 \leq i \leq \ell$ by construction, it also holds that $\nu'(\text{body}(\gamma)) \leq \inf_{\nu \in S} \nu(\text{body}(\gamma))$.

Therefore also $\nu'(\text{body}(\gamma)) \leq (1 - K) + \inf_{\nu \in S} \nu(\text{head}(\gamma))$, since all $\nu \in S$ are K -fuzzy models. That is, there exists a rational $r = \max\{0, \nu'(\text{body}(\gamma)) - 1 + K\}$ (recall that we only consider rational K) such that r is less or equal to $\nu(\text{head}(\gamma))$ for every $\nu \in S$ and $\nu'(\text{body}(\gamma)) \leq (1 - K) + r$. Then, $\nu'(\text{head}(\gamma)) = \inf_{\nu \in S} \nu(\text{head}(\gamma)) \geq r$ and therefore also $\nu'(\text{body}(\gamma)) \leq (1 - K) + \nu'(\text{head}(\gamma))$, i.e., ν' K -satisfies γ . We see that ν' is indeed a K -fuzzy model and clearly for every $\nu \in S$ we also have $\nu' \leq \nu$. Taking the infimum over all K -fuzzy models of (Π, τ) then yields the unique minimal K -fuzzy model for (Π, τ) . \square

Proposition 2. Let G be a ground atom, $c \in [0, 1]$, and μ the minimal K -fuzzy model of (Π, τ) . Then $\mu(G) \geq c$ if and only if $\nu(G) \geq c$ for all K -fuzzy models of (Π, τ) .

¹ Considering $\nu(G) \leq \tau(G)$ allows the trivial model setting everything to truth 0

Of particular importance to our main result is the fact that fuzzy models are, in a sense, only more fine-grained versions of classical models. In particular, since the classical case corresponds to every observation having maximal truth degree, fuzzy models are in a sense bounded from above by classical models.

Lemma 3. *Let (Π, τ) be an MV-instance. Let ν be a minimal K -fuzzy model of (Π, τ) . Then for every $G \in GAtoms$, we have $\nu(G) > 0$ only if $\Pi^{crisp}, D_\tau \models G$.*

3.2 Characterising Minimal Fuzzy Models

Minimal K -fuzzy models are a natural tool for query answering. In the following we show that we can in fact characterise the minimal K -fuzzy model for (Π, τ) in terms of a simple linear program $Opt_{\Pi, \tau}^K$ defined over the ground rules induced by the oblivious chase sequence² for Π^{crisp}, D_τ .

Let (Π, τ) be an MV-instance. Let $\Gamma = \{\gamma_1, \dots, \gamma_m\}$ be the ground rules in $OGround(\Pi^{crisp}, D_\tau)$. Let $\mathcal{G} = \{G_1, \dots, G_n\}$ be all of the ground atoms occurring in rules in Γ . For every $G_i \in \mathcal{G}$ we associate G_i with a variable x_i in our linear program that intuitively will represent the truth degree of G_i . For γ_j of the form $G_{j_1}, \dots, G_{j_\ell} \rightarrow G_{j_{head}}$ define

$$\text{Eval}(\gamma_j) := \sum_{k=1}^{\ell} (1 - x_{j_k}) + x_{j_{head}}$$

which directly expresses the satisfaction of rule γ_j , with variable x_{j_k} representing the truth of G_{j_k} . The linear program $Opt_{\Pi, \tau}^K$ is then defined as follows

$$\begin{aligned} & \text{minimise} && \sum_{i=1}^n x_i \\ & \text{subject to} && \text{Eval}(\gamma_j) \geq K && \text{for } 1 \leq j \leq m \\ & && x_i = \tau(G_i) && \text{for } i \text{ where } \tau(G_i) \text{ is defined} \\ & && 1 \geq x_i \geq 0 && \text{for } 1 \leq i \leq n \end{aligned} \tag{1}$$

Recall, a *solution* of a linear program is any assignment to the variables, a *feasible solution* is a solution that satisfies all constraints, and an *optimal solution* is a feasible solution with minimal value of the objective function. With this in mind we can state our main result for MV-Datalog, a characterisation of minimal K -fuzzy models in terms of optimal solutions of $Opt_{\Pi, \tau}^K$.

Theorem 4. *Let Π, τ be an MV-instance. Then \mathbf{x} is an optimal solution of $Opt_{\Pi, \tau}^K$ if and only if $\nu_{\mathbf{x}}$ is a minimal K -fuzzy model of (Π, τ)*

By construction, any feasible solutions of $Opt_{\Pi, \tau}^K$ induces a K -fuzzy model $\nu_{\mathbf{x}}$ that assigns $\nu_{\mathbf{x}}(G_i) = x_i$ and 0 to all other ground atoms not in \mathcal{G} . Similarly, in the other direction we can observe that any K -fuzzy model ν of (Π, τ) can be used to construct a feasible solution of $Opt_{\Pi, \tau}^K$. This natural correspondence of solutions of $Opt_{\Pi, \tau}^K$ and K -fuzzy models of (Π, τ) is formalised in the following lemma.

Lemma 5. *Let Π, τ be an MV-instance. The following statements hold:*

² While the use of the oblivious chase without existential quantification is untypical we explicitly require all rules induced by oblivious applications for our characterisation.

1. For any feasible solution \mathbf{x} of $Opt_{\Pi, \tau}^K$, $\nu_{\mathbf{x}}$ is a K -fuzzy model of (Π, τ) .
2. For any K -fuzzy model ν of (Π, τ) , the solution \mathbf{x} with $x_i = \nu(G_i)$ is a feasible solution of $Opt_{\Pi, \tau}^K$.
3. $Opt_{\Pi, \tau}^K$ is feasible if and only if (Π, τ) is K -satisfiable.

We say that a ground rule $\gamma \in \Gamma$ is ν -tight if $\nu(\gamma) = K$, that is $\nu(\text{body}(\gamma)) - 1 + K = \nu(\text{head}(\gamma))$. When a rule is not tight we refer to $\nu(\text{head}(\gamma)) - (\nu(\text{body}(\gamma)) - 1 + K)$ as the $\nu_{\mathbf{x}}$ -gap of the rule. The gap is at least 0 for any K -satisfied ground rule. We will be interested in the structural relationships between tight rules. We refer to the ground atoms in \mathcal{G} for which τ is not defined as the *derived ground atoms* of $Opt_{\Pi, \tau}^K$.

Lemma 6. *Let ν be a minimal model of an MV-instance Π, τ . Let \mathcal{G}' be a non-empty set of ground atoms for which ν assigns truth greater than 0. Then there exists at least one ν -tight ground rule in $OGround(\Pi^{crisp}, D_{\tau})$ such that $\text{head}(\gamma) \in \mathcal{G}'$ and $\text{body}(\gamma) \cap \mathcal{G}' = \emptyset$.*

Proof of Theorem 4. Observe, that by Lemma 3 and the fact that the chase produces all atoms that can be entailed by (Π, τ) , it follows that we do not have to consider solutions with atoms other than the set \mathcal{G} considered in the definition of $Opt_{\Pi, \tau}^K$.

For the if direction, let μ be a minimal model of (Π, τ) with $\mu(G) = 0$ for all ground atoms $G \notin \mathcal{G}$. By Lemma 5, μ induces a feasible solution \mathbf{x} of $Opt_{\Pi, \tau}^K$ where $x_i = \mu(G_i)$. Suppose this solution is not optimal and let \mathbf{y} be an optimal solution of $Opt_{\Pi, \tau}^K$. Let $\mathcal{G}' \subseteq \mathcal{G}$ be the set of ground atoms G where $\nu_{\mathbf{y}}(G) < \mu(G)$. Since we assume that $\sum y_i < \sum x_i$, this set is not empty. By Lemma 6 there exists a μ -tight rule γ where the head is in \mathcal{G}' but none of the body atoms of γ are. It follows that

$$\nu_{\mathbf{y}}(\text{head}(\gamma)) < \mu(\text{head}(\gamma)) = \mu(\text{body}(\gamma)) - 1 + K = \nu_{\mathbf{y}}(\text{body}(\gamma)) - 1 + K$$

Hence, $\nu_{\mathbf{y}}$ does not K -satisfy γ . By Lemma 5 this contradicts the assumption that \mathbf{y} is a feasible solution with lower objective than \mathbf{x} .

For the other direction, let \mathbf{x} be an optimal solution of $Opt_{\Pi, \tau}^K$. Suppose now that $\nu_{\mathbf{x}}$ is not a minimal K -fuzzy model of (Π, τ) . Hence, there exists some model ν' with $\nu' < \nu_{\mathbf{x}}$. Let \mathbf{y} be the solution of $Opt_{\Pi, \tau}^K$ induced by ν' . Since at least one atom is less true, and none are more true in ν' we have $\sum y_i < \sum x_i$, contradicting the optimality of \mathbf{x} for $Opt_{\Pi, \tau}^K$. \square

Since the oblivious chase is polynomial with respect to data complexity (without existential quantification) and finding optimal solutions of linear programs is famously polynomial Khachiyan (1979), our characterisation of minimal K -fuzzy models also directly reveals the complexity of K -TRUTH. Recall that for $K = 1$, \mathbf{P} -hardness is inherited from classical Datalog (see Dantsin et al. (2001)).

Corollary 7. *Fix a rational $K \in (0, 1]$. K -TRUTH is in \mathbf{P} with respect to data complexity. Moreover, 1-TRUTH is \mathbf{P} -complete in data complexity.*

To conclude this section we briefly note why it is important to use the oblivious chase as a basis for the ground rules that make up $Opt_{\Pi, \tau}^K$. Consider the program $P(x), Q(x) \rightarrow S(x)$ and $R(x) \rightarrow S(x)$. In plain Datalog, $S(a)$ can be inferred simply if either $P(a)$ and $Q(a)$, or $R(a)$ are in the database. However, in the many-valued setting the truth of $S(a)$ depends on the truth of all three facts, $P(a)$, $Q(a)$, and $R(a)$ since it needs to be at least as true as $P(x) \odot Q(a)$ and at least as true $R(a)$. Technically this means $Opt_{\Pi, \tau}^K$ needs to

consider all possible paths of inferring each ground atom, and this exactly corresponds to the oblivious chase applying rules once per homomorphism from the body into the instance.

3.3 Classical Behaviour for Certain Knowledge

We are particularly interested in settings where one wants to combine large amounts of certain knowledge, e.g., classical knowledge graphs and databases, with additional fuzzy information. Going back to our example from the introduction, we have such a scenario if we want to reason on properties of the animals in the pictures by using pre-existing biological and zoological ontologies. Ideally, we want inference on the ontologies to behave exactly as in the Datalog and only differ in behaviour when such *certain knowledge* is combined with fuzzy observations and consequences of fuzzy observations. Note that particularly in the context of ontologies the question of fuzziness has also been studied extensively in the field of fuzzy description logics, see e.g., Borgwardt and Peñaloza (2017); Stoilos et al. (2007); Lukasiewicz and Straccia (2008). Notably, the \mathcal{EL} family of description logics, which is closely related to rule based languages, has also been studied specifically with fuzzy semantics based on Lukasiewicz logic and shown undecidable by Borgwardt et al. (2017).

We show now that MV-Datalog indeed exhibits such ideal behaviour. In particular, when we consider Lemma 3 but with a database consisting only of certain knowledge in τ , i.e, ground atoms for which τ evaluates to 1, we can state the following.

Theorem 8. *Let Π, τ be a MV-Datalog program and database. Let D^1 be the classical database containing all $G \in GAtoms$ where $\tau(G) = 1$. If Π, τ is 1-satisfiable, then for every ground atom G , $\Pi^{crisp}, D^1 \models G$ if and only if $\mu(G) = 1$ where μ is the minimal 1-fuzzy model of (Π, τ) .*

Beyond the conceptual importance of Theorem 8, the statement also has significant practical consequences. Our characterisation of minimal K -fuzzy models provides a clear algorithm for computing minimal K -fuzzy models via the oblivious chase and then solving $Opt_{\Pi, \tau}^K$. Using the oblivious chase is important for our characterisation, since different ground rules with the same head may differ in how true their body is. In classical Datalog it is of course sufficient to infer the head once, which can require significantly less effort in practice. According to Theorem 8 (and by uniqueness of minimal fuzzy models) it is not necessary to compute the truth degree of those facts that are entailed by Π^{crisp}, D^1 via $Opt_{\Pi, \tau}^K$. An implementation can use this observation to produce less ground rules (i.e., not the full oblivious chase) and smaller equivalent versions of $Opt_{\Pi, \tau}^K$ with less computational effort. We conclude that MV-Datalog is well suited for combining large amounts of certain knowledge with fuzzy observations. Not only are the semantics on certain data equivalent to those of classical Datalog, but the computation of minimal fuzzy models also exhibits desirable properties in the presence of certain data.

4 MV-Datalog[±]

Existential quantification in the head of rules provide a natural way of dealing with the common problem of incomplete or missing data by declaratively stating that certain facts

must exist. Such usage is especially powerful in applications where rule-based reasoning is used for complex data analysis (see e.g., Bellomarini et al. (2017)). In the following, we introduce MV-Datalog[±] as an extension of MV-Datalog with a focus on providing a fuzzy reasoning language that is useful for such applications.

To obtain semantics matching our intuition for such a system, we propose an alternative to the commonly studied semantics for existential quantification for Łukasiewicz logic. We then identify certain *preferred models* that exhibit desired conceptual and computational behaviour as the basis for reasoning in MV-Datalog[±].

4.1 Strong Existential Quantification in Łukasiewicz Semantics

The semantics of existential quantification in Łukasiewicz logic is commonly defined as $\nu(\exists \mathbf{x} \varphi(\mathbf{x})) = \sup\{\nu(\varphi[\mathbf{x}/\mathbf{c}]) \mid \mathbf{c} \in \text{Dom}^{|\mathbf{x}|}\}$ where $\varphi[\mathbf{x}/\mathbf{c}]$ is the substitution of variables \mathbf{x} in φ by constants \mathbf{c} (cf. Hájek (1998)).

However, with our focus set on practical applications, we propose alternative semantics for existential quantification that match the semantics of \oplus . To contrast between the aforementioned semantics of \exists we refer to our semantics as *strong existential quantification*.

$$\nu(\exists \mathbf{x} \varphi(\mathbf{x})) = \min\{1, \sum_{\mathbf{c} \in \text{Dom}^{|\mathbf{x}|}} \nu(\varphi[\mathbf{x}/\mathbf{c}])\}$$

We refer to \mathbf{L} extended with strong existential quantification as \mathbf{L}_{\exists} (following standard syntax of first order existential quantification). Such semantics for quantification in Łukasiewicz logic have also recently been studied more generally in a purely logical context by Fjellstad and Olsen (2021).

Definition 3. A MV-Datalog[±] program is a set of \mathbf{L}_{\exists} formulas that are either MV-Datalog rules or of the form

$$B_1 \odot B_2 \odot \dots \odot B_n \rightarrow \exists \mathbf{x} \varphi(\mathbf{x}, \mathbf{y})$$

where B_i are relational atoms, \mathbf{y} are the variables in the body, and φ is formula that contains only a relational atom that uses all variables of \mathbf{x} .

Note that we generally do not restrict the domain in the semantics of \mathbf{L} formulas. Thus, as in Datalog[±], interpretations are not constrained to the active domain of Π and τ but allow for the introduction of new constants. As noted above, the introduction of new constants provides a natural mechanism to deductively handle missing and incomplete data. Our proposed use of strong existential quantification is motivated by its behaviour in cases of incomplete information as illustrated in the following example.

Example 2. We consider the following example rule, expressing that every company has a key person, to illustrate the difference between the two semantics for \exists .

$$\text{Company}(y) \rightarrow \exists x \text{ KeyPerson}(x, y)$$

Suppose we have the following database, we know for certain that Acme is a company and we are 0.8 degrees confident that Amy is the key person of Acme.

$$\tau(\text{KeyPerson}(\text{Amy}, \text{Acme})) = 0.8, \quad \tau(\text{Company}(\text{Acme})) = 1$$

We discuss the case in a 1-fuzzy model in the following. With existential semantics via the supremum of matching ground atoms, the first rule implies the existence of some new $\text{KeyPerson}(\text{N}_1, \text{Acme})$ with truth degree 1. With strong existential quantification, the first rule implies only truth degree 0.2 for some $\text{KeyPerson}(\text{N}_1, \text{Acme})$ since the known observation $\text{KeyPerson}(\text{Amy}, \text{Acme})$ already contributes 0.8 degrees of truth to the head.

Intuitively, the new constant N_1 assumes the role of the unknown other object that possibly is the key person. In traditional semantics for \exists we have to infer (in a 1-fuzzy model) that N_1 is certain to control Acme . In contrast, under strong existential semantics, the confidence in N_1 being the key person is determined by the known observation that Amy might be the key person.

4.2 Preferred Models for MV-Datalog[±]

When considered in full generality, our proposed strong existential quantification semantics reveal a variety of theoretical challenges. To allow for K -fuzzy models ν where $\{G \in \text{GAtoms} \mid \nu(G) > 0\}$ is not finite would require the introduction of further technical complexity in the definition. Furthermore, it can be necessary for a model to introduce multiple new constants per ground rule to satisfy an existential quantification in certain situations as illustrated by the following example.

Example 3. Consider database τ with $\tau(S(a)) = 0.8, \tau(T(a)) = 0.2$ and program Π

$$\begin{aligned} S(x) &\rightarrow \exists y P(x, y) \\ P(x, y) &\rightarrow T(x) \end{aligned}$$

Then there is no 1-fuzzy model with only one new constant. However, there is a solution where the facts $P(a, n_1), P(a, n_2), P(a, n_3), P(a, n_4)$, with new constants n_1, \dots, n_4 , are all assigned truth degree 0.2.

Infinite models and the phenomenon illustrated by Example 3 present interesting topics for further theoretical study. For practical applications we propose to avoid these corner cases and focus on natural *preferred models* that avoid these issues and provide predictable outcomes.

Recall that the oblivious chase sequence introduces new constants exactly once for every possible true grounding of the body. This limitation on the number of new constants provides a natural balance between not overly restricting facts with null values in models while also avoiding extreme situations as in Example 3 where arbitrarily many constants are introduced to satisfy the rule for a single ground body. We also restrict minimality of fuzzy models to only those ground atoms without labelled nulls. Minimising also the truth of atoms with labelled nulls would intuitively correspond to a kind of soft closed-world assumption, where we want to infer truth of statements under the assumption that the unknown information is minimal.

Definition 4. Let Π, τ be an MV^\pm -instance. Let μ, ν be two K -fuzzy models of (Π, τ) . We say that $\mu \leq_a \nu$ if for all $G \in \text{GAtoms}[\text{Adom}]$ it holds that $\mu(G) \leq \nu(G)$.

Definition 5. We say that a K -fuzzy model ν of MV-Datalog^\pm instance (Π, τ) has an oblivious base if $\text{OLim}(\Pi^{\text{crisp}}, D_\tau)$ is finite, if for every $G \in \text{GAtoms}$ it holds that $\nu(G) > 0$ only if $G \in \text{OLim}(\Pi^{\text{crisp}}, D_\tau)$. We say that ν is a preferred K -fuzzy model if it

is minimal with respect to the partial ordering \leq_a of all K -fuzzy models with an oblivious base .

This definition of preferred models captures the previously outlined intuition of which kinds of models we want to consider for practical applications. That is, we propose preferred models as the desired output of an MV-Datalog $^\pm$ system. An MV $^\pm$ -instance does not necessarily have a unique preferred model. The question of how to compare preferred models among each other is left as an open question.

We now extend the definition of $Opt_{\Pi,\tau}^K$ from Section 3.1 to program $\exists\text{-}Opt_{\Pi,\tau}^K$ for an MV $^\pm$ -instance Π, τ . Let \mathcal{G} and Γ be as in the previous definition of $Opt_{\Pi,\tau}^K$. For ground rules that are the result of grounding a rule without existential quantification, Eval is defined the same as it was for $Opt_{\Pi,\tau}^K$. Otherwise, let γ be a ground rule obtained from grounding of an existential rule. Let G_h be the atom in the head of γ and let N be the labelled nulls at the positions in G_h that were existentially quantified in the original rule before grounding. We say that a ground atom G matches G_h if G can be obtained from G_h by replacing the labelled nulls in N by some other values of the domain. Let $G_{\beta_1}, \dots, G_{\beta_k}$ be the ground atoms in \mathcal{G} that match G_h . We can then extend our previous definition of Eval to groundings of existential rules as

$$\text{Eval}(\gamma) := \sum_{j=1}^{\ell} (1 - x_{\alpha_j}) + \sum_{i=1}^k x_{\beta_i}$$

with variables x_{α_j} and x_{β_i} corresponding, respectively, to the truth values of the ground atoms in the body and the head. Then $\exists\text{-}Opt_{\Pi,\tau}^K$ is the linear program defined the same as $Opt_{\Pi,\tau}^K$, but with objective function

$$\text{minimise } \sum_{i=1}^n w_i x_i \tag{2}$$

where x_i provides the truth value of G_i , $w_i = 1$ if $G_i \in GAtoms[Adom]$ and $w_i = 0$ otherwise. That is, ground atoms that contain nulls do not contribute to the objective. Recall from Example 3 that this does not imply that they can be simply set to be fully true.

Theorem 9. *Let Π, τ be a MV-Datalog program and database where $OLim(\Pi^{crisp}, D_\tau)$ is finite. The following two statements hold.*

1. *If (Π, τ) has a preferred K -fuzzy model, then $\exists\text{-}Opt_{\Pi,\tau}^K$ is feasible.*
2. *For any optimal solution \mathbf{x} of $\exists\text{-}Opt_{\Pi,\tau}^K$, we have that $\nu_{\mathbf{x}}$ is a preferred K -fuzzy model for (Π, τ) .*

From our definition of preferred models we are naturally interested in identifying cases where the oblivious chase is finite. There exist syntactic restrictions for which we can show that this is always the case. In particular, we can identify programs where the oblivious chase is finite based on the notion of weakly acyclic programs as studied by Fagin et al. (2005). Technical details of this question are discussed in Appendix B.

5 Conclusion & Outlook

We have introduced MV-Datalog and MV-Datalog $^\pm$, two many-valued rule-based reasoning languages following Łukasiewicz semantics. We provide theoretical foundations for

MV-Datalog by characterising its minimal models in terms of an optimisation problem and showing that fuzzy fact entailment in MV-Datalog has the same complexity as fact entailment in Datalog. For MV-Datalog[±] we introduce alternative semantics for existential quantification in Łukasiewicz logic that better match the popular usage of \exists for handling missing data. As a first starting point for the practical use of MV-Datalog[±] we introduce preferred models and show how to compute them.

We argue that MV-Datalog and its extensions can be effective languages for reasoning in settings with uncertain data. Therefore the development of a system based on the results given here is the natural next step to validate efficiency and effectiveness in comparison to other formalisms. This also entails the study of how the computation of minimal or preferred models can be further optimised, ideally requiring fewer ground rules than the methods described here.

Beyond practical considerations, a variety of natural theoretical questions have been left open. Prime among them are the questions of how fuzzy models of MV-Datalog behave with stratified negation, and how to handle the behaviour of MV-Datalog[±] instances with an infinite oblivious chase. The latter is particularly intriguing since truth degrees effectively decreases monotonically along paths of derivation. We therefore suspect there to be relevant classes of problems where the oblivious chase is infinite but instances always permit finite fuzzy models.

Finally, there is an intriguing possibility of combining MV-Datalog semantics with the probabilistic semantics of Problog to reason about uncertainty and likelihood as two orthogonal concepts at the same time.

Acknowledgements

Stefano Sferrazza was supported by the Austrian Science Fund (FWF):P30930. Georg Gottlob is a Royal Society Research Professor and acknowledges support by the Royal Society in this role through the “RAISON DATA” project (Reference No. RP\R1\201074). Matthias Lanzinger acknowledges support by the Royal Society “RAISON DATA” project (Reference No. RP\R1\201074).

References

- ACHS, Á. AND KISS, A. 1995. Fuzzy extension of datalog. *Acta Cybern.*, 12, 2, pp. 153–166.
- BACH, S. H., BROECHELER, M., HUANG, B., AND GETOOR, L. 2017. Hinge-Loss Markov Random Fields and Probabilistic Soft Logic. *J. Mach. Learn. Res.*, 18, pp. 109:1–109:67.
- BELLOMARINI, L., GOTTLÖB, G., PIERIS, A., AND SALLINGER, E. Swift logic for big data and knowledge graphs. In *Proc. IJCAI 2017*, pp. 2–10. ijcai.org.
- BELLOMARINI, L., LAURENZA, E., SALLINGER, E., AND SHERKHONOV, E. Reasoning Under Uncertainty in Knowledge Graphs. In *RuleML+RR 2020*, pp. 131–139. Springer.
- BOBILLO, F. AND STRACCIA, U. 2016. The fuzzy ontology reasoner fuzzydl. *Knowl. Based Syst.*, 95, pp. 12–34.
- BORGWARDT, S., CERAMI, M., AND PEÑALOZA, R. Łukasiewicz fuzzy EL is undecidable. In *Proc. DL 2017*, volume 1879 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- BORGWARDT, S. AND PEÑALOZA, R. Fuzzy description logics - A survey. In *Proc. SUM 2017*, volume 10564 of *Lecture Notes in Computer Science*, pp. 31–45. Springer.
- CALÌ, A., GOTTLÖB, G., AND KIFER, M. 2013. Taming the infinite chase: Query answering under expressive relational constraints. *JAIR*, 48, pp. 115–174.

- CALÌ, A., GOTTLOB, G., AND LUKASIEWICZ, T. 2012. A general Datalog-based framework for tractable query answering over ontologies. *Journal of Web Semantics*, 14, pp. 57–83.
- CAVALLO, R. AND PITTARELLI, M. The theory of probabilistic databases. In *VLDB’87* 1987, pp. 71–81. Morgan Kaufmann.
- DANTSIN, E., EITER, T., GOTTLOB, G., AND VORONKOV, A. 2001. Complexity and expressive power of logic programming. *ACM Comput. Surv.*, 33, 3, pp. 374–425.
- EBRAHIM, R. 2001. Fuzzy logic programming. *Fuzzy Sets Syst.*, 117, 2, pp. 215–230.
- FAGIN, R., KOLAITIS, P. G., MILLER, R. J., AND POPA, L. 2005. Data exchange: semantics and query answering. *Theor. Comput. Sci.*, 336, 1, pp. 89–124.
- FJELLSTAD, A. AND OLSEN, J.-F. 2021. IKT^w and Lukasiewicz-Models. *Notre Dame Journal of Formal Logic*, 62, 2, pp. 247 – 256.
- HÁJEK, P. 1998. *Metamathematics of Fuzzy Logic*, volume 4 of *Trends in Logic*. Kluwer.
- IRANZO, P. J. AND SÁENZ-PÉREZ, F. 2018. A fuzzy datalog deductive database system. *IEEE Trans. Fuzzy Syst.*, 26, 5, pp. 2634–2648.
- KHACHIYAN, L. G. A polynomial algorithm in linear programming. In *Doklady Akademii Nauk* 1979, volume 244, pp. 1093–1096. Russian Academy of Sciences.
- LEE, H. S., JUNG, H., AGARWAL, A. A., AND KIM, J. 2017. Can deep neural networks match the related objects?: A survey on imagenet-trained classification models. *CoRR*, abs/1709.03806.
- LUKASIEWICZ, T. AND STRACCIA, U. 2008. Managing uncertainty and vagueness in description logics for the semantic web. *J. Web Semant.*, 6, 4, pp. 291–308.
- MAIER, D., TEKLE, K. T., KIFER, M., AND WARREN, D. S. Datalog: concepts, history, and outlook. In *Declarative Logic Programming: Theory, Systems, and Applications* 2018. ACM / Morgan & Claypool.
- MISHKIN, D., SERGIEVSKIY, N., AND MATAS, J. 2017. Systematic evaluation of convolution neural network advances on the imagenet. *Comput. Vis. Image Underst.*, 161, pp. 11–19.
- PREINING, N. Gödel logics - A survey. In *Proc. LPAR* 2010, volume 6397 of *Lecture Notes in Computer Science*, pp. 30–51. Springer.
- RAEDT, L. D., KIMMIG, A., AND TOIVONEN, H. Problog: A probabilistic prolog and its application in link discovery. In *Proc. IJCAI* 2007, pp. 2462–2467.
- RICHARDSON, M. AND DOMINGOS, P. M. 2006. Markov logic networks. *Mach. Learn.*, 62, 1-2, pp. 107–136.
- STOILLOS, G., STAMOU, G. B., PAN, J. Z., TZOUVARAS, V., AND HORROCKS, I. 2007. Reasoning with very expressive fuzzy description logics. *J. Artif. Intell. Res.*, 30, pp. 273–320.
- SUCIU, D., OLTEANU, D., RÉ, C., AND KOCH, C. 2011. *Probabilistic Databases*. Synthesis Lectures on Data Management. Morgan & Claypool Publishers.
- VOJTÁS, P. 2001. Fuzzy logic programming. *Fuzzy Sets Syst.*, 124, 3, pp. 361–370.

Appendix A Proof Details

Proof of Lemma 3. Let G be a ground atom with $\nu(G) > 0$. Since ν is minimal there must be a ground rule with G in the head such that for every body atom $G_i \in \text{body}(\gamma)$ also $\nu(G_i) > 0$. Consider the tree with γ as the root and a child for every body atom of γ . At each child take again a rule where the respective body atom is in the head and the body is true to some degree greater than 0.

Repeating the process, we must arrive at a finite tree where all leaves must correspond to ground rules where all body atoms are defined in τ with truth greater than 0. Otherwise, we could simply set $\nu(G') = 0$ for all ground atoms G' whose branches do not lead to such leaves and ν would still be a K -fuzzy model, contradicting its minimality.

Viewing the same tree in the context of $\Pi^{\text{crisp}}, D_\tau$ we then see, going from the leaves upward that the bodies of all these rules must be true in all models. Thus, also ultimately $\Pi^{\text{crisp}}, D_\tau \models G$. \square

Proof of Lemma 5. Let \mathbf{x} be a feasible solution of $\text{Opt}_{\Pi, \tau}^K$. Observe that $\text{Eval}(\gamma) = \nu_{\mathbf{x}}(\gamma)$ for every ground rule γ : Since, $\text{Eval}(\gamma) \geq K$ for all $\gamma \in \Gamma$, we also see that $\nu_{\mathbf{x}}$ K -satisfies every rule in $\text{OGround}(\Pi^{\text{crisp}}, D_\tau)$. By definition of an oblivious chase sequence and since $\nu_{\mathbf{x}}(G) = 0$ for any ground atom that is not mentioned in $\text{OLim}(\Pi^{\text{crisp}}, D_\tau)$, any other ground rule $\gamma \notin \Gamma$ is trivially K -satisfied since both head and body have truth 0.

Let ν be a K -fuzzy model of (Π, τ) . Let \mathbf{x} be a solution of $\text{Opt}_{\Pi, \tau}^K$ with $x_i = \nu(G_i)$ for all $G \in \mathcal{G}$. Since ν K -satisfies all ground rules, we in particular have that $\nu(\gamma) = \text{Eval}(\gamma) \geq K$ for all $\gamma \in \Gamma$ and therefore \mathbf{x} is feasible. The third statement follows immediately from combination of the first two. \square

Proof of Lemma 6. By Lemma 3, and the fact that the oblivious chase sequence constructs minimal models for Datalog, we also have that $\mathcal{G}' \subseteq \text{OLim}(\Pi^{\text{crisp}}, D_\tau)$. It follows by construction that every derived $G \in \mathcal{G}$ in the head of some ground rule in $\text{OGround}(\Pi^{\text{crisp}}, D_\tau)$. What is left to show is that at least one such rule is also ν -tight.

Suppose the statement is false and let $\delta > 0$ be the minimal ν -gap of rules in Γ whose body atoms are all not in \mathcal{G}' . Let Γ' contain only the ν -tight rules in $\text{OGround}(\Pi^{\text{crisp}}, D_\tau)$ that contain an atom from \mathcal{G}' in the head. Since we assume the statement false, all of the rules in Γ' have a body atom that is also in \mathcal{G}' . Now consider the truth assignment ν' defined as

$$\nu'(G) = \begin{cases} \max\{0, \nu(G) - \delta\} & \text{if } G \in \mathcal{G}' \\ \nu_{\mathbf{x}}(G) & \text{otherwise} \end{cases}$$

By assumption we have that all ν -tight rules where \mathcal{G}' occurs in the head also have some atom from \mathcal{G}' in the body. Meaning that under ν' the truth of both head and body are decreased by at least δ (or until both are 0). Hence, all rules in Γ' are still K -satisfied. We can also see that all other rules that contain an atom in \mathcal{G}' remain K -satisfied under ν' . When an atom of \mathcal{G}' occurs in the body of a rule, the implication can only become more true. Where an atom G from \mathcal{G}' occurs only in the head of a rule $\gamma \notin \Gamma'$, we have that the $\nu_{\mathbf{x}}$ -gap of the rule is at least δ . That is,

$$\nu_{\mathbf{x}}(G) \geq \nu_{\mathbf{x}}(\text{body}(\gamma)) - 1 + K + \delta$$

and therefore also

$$\nu'(G) \geq \nu_{\mathbf{x}}(G) - \delta \geq \nu_{\mathbf{x}}(\text{body}(\gamma)) - 1 + K \geq \nu'(\text{body}(\gamma)) - 1 + K$$

We therefore see that ν' is a K -fuzzy model with $\nu' < \nu$ and we arrive at a contradiction. \square

Proof of Theorem 8. The implication from right to is a special case of Lemma 3. For the other direction, we argue by induction over the oblivious chase sequence D_0, D_1, \dots for Π^{crisp}, D^1 , that if some ground atom G is in D_i , then $\mu(G) = 1$ in every 1-fuzzy model μ of (Π, τ) . The base case $D_0 = D^1$ follows by definition of D^1 . For the induction, suppose the claim holds up to step i . In the step from D_i to D_{i+1} either the ground atoms in both are the same (the oblivious application produced a ground head that was already in D_i) or there is a single new ground atom G in D_{i+1} but not in D_i . Let γ be the ground rule induced by the oblivious application in the step from i to $i + 1$. Since the rule was applicable, for all $G' \in \text{body}(\gamma)$ we have $G' \in D_i$ and by the induction hypothesis also $\mu(G') = 1$. Hence, $\mu(\text{body}(\gamma)) = 1$ and since γ must be 1-satisfied, we have that in every 1-fuzzy model also $\mu(G) = 1$.

Recall, that $OLim(\Pi^{crisp}, D^1)$ is always a minimal model of Π^{crisp}, D^1 and thus $\Pi^{crisp}, D^1 \models G$ if and only if $G \in OLim(\Pi^{crisp}, D^1)$. And by the above induction argument $G \in OLim(\Pi^{crisp}, D^1)$ implies that $\mu(G) = 1$. \square

Proof of Theorem 9. First, observe that since $OLim(\Pi^{crisp}, D_\tau)$ is unique up to isomorphism, we can assume without loss of generality that there is no preferred model where some ground atom G is true but not in the set \mathcal{G} of ground atoms considered in the construction of $\exists\text{-}Opt_{\Pi, \tau}^K$.

Let ν be a preferred K -fuzzy model of Π, τ . Since ν has an oblivious base, we have that $\nu(G) > 0$ only if $G \in \mathcal{G}$. Let \mathbf{x} be the solution of $\exists\text{-}Opt_{\Pi, \tau}^K$ where $x_i = \nu(G_i)$ for all $G_i \in \mathcal{G}$. Then, by definition of Eval we immediately see $\text{Eval}(\gamma) = \nu(\gamma)$ for every ground rule γ . Hence, if all ground rules are K -satisfied by ν , then all constraints $\text{Eval}(\gamma) \geq K$ are satisfied and ν is feasible.

For the second statement, assume that \mathbf{x} is an optimal solution of $\exists\text{-}Opt_{\Pi, \tau}^K$ but $\nu_{\mathbf{x}}$ is not preferred. The only way $\nu_{\mathbf{x}}$ can not be preferred is if it is not active minimal, i.e., there is some K -fuzzy model μ where for all $G \in GAtoms[Adom]$, $\mu(G) \leq \nu_{\mathbf{x}}(G)$ and for at least one $G' \in GAtoms[Adom]$, $\mu(G') < \nu_{\mathbf{x}}(G')$. Since both μ and $\nu_{\mathbf{x}}$ have an oblivious base (which is unique up to isomorphism) we see that every $G \in GAtoms$ where $\mu(G) > 0$ is also in \mathcal{G} . Hence, it is straightforward to construct a feasible solution \mathbf{x}' from μ for which the objective function is strictly lower than for \mathbf{x} , a contradiction. \square

Appendix B Finiteness in the Oblivious Chase

Since we are interested particularly in instances where the chase is finite, it is of interest to identify fragments where this is always the case. The most prominent condition for which the finiteness of the chase in Datalog[±] is studied, is the restriction to *weakly acyclic programs* as first introduced by Fagin et al. (2005). Let Π be a set of rules with over the relational language σ . We first define the *dependency graph* for Π as the graph with vertices $\{(R, i) \mid R \in \sigma, i \in \{1, \dots, \#(R)\}\}$. We say a variable x is in position (R, i) if x is at the i -th place of a relation with name R . There is a (normal) edge from vertex (R, i) to

(S, j) exactly if there is a rule in Π where the same variable x occurs at position (R, i) in the body and at position (S, j) in the head of the rule. There is a *special edge* from (R, i) to (S, j) exactly if there is an existential rule in Π , such that there is a variable x that x occurs in both the body and head, x is in position (R, i) of the body, and an existentially quantified y is in position (S, j) in the head. We say that a program is *weakly acyclic* if its dependency graph has no cycle that passes a special edge.

However, the standard definition of weakly acyclic is particular to the restricted chase, which does not fit our setting. This becomes apparent when considering the simple (weakly acyclic) program

$$P(x) \rightarrow \exists y P(y).$$

The oblivious chase is infinite in this case, since every new instantiation of y generates a new homomorphism that satisfies the body. Since x does not occur in the head at all, the dependency graph used in the standard definition of weakly-acyclic does not have any edges and is thus weakly acyclic.

Let the *variable expansion* $\text{ve}(\Pi)$ of Π be the program obtained by the following rewriting. Without loss of generality every existential rule is of the form $\varphi(\mathbf{x}) \rightarrow \exists \mathbf{y} R(\mathbf{y}, \mathbf{x}')$, where φ is the body formula, R is a relation symbol and \mathbf{x}' are those (free) variables that occur in both the body and the head. Let \mathbf{x}^* be the free variables that only occur in φ but not in the head. Then, for every such rule we replace it by the two rules

$$\begin{aligned} \varphi(\mathbf{x}) &\rightarrow \exists \mathbf{y} R^*(\mathbf{y}, \mathbf{x}', \mathbf{x}^*) \\ R^*(\mathbf{y}, \mathbf{x}', \mathbf{x}^*) &\rightarrow R(\mathbf{y}, \mathbf{x}') \end{aligned}$$

to obtain $\text{ve}(\Pi)$. Intuitively, the variable expansion reveals those particular cycles that are harmless in the restricted chase but lead to infinite sequences in the oblivious chase in the standard definition of weak acyclicity. For example, the variable expansion of our simple example above would thus be the (no longer weakly acyclic) program

$$\begin{aligned} P(x) &\rightarrow \exists y P^*(y, x) \\ P^*(y, x) &\rightarrow P(y) \end{aligned}$$

The following theorem then follows by similar argument as originally given for the restricted chase (Fagin et al. 2005, Theorem 3.9) by additionally observing that nulls can never propagate to a position in the dependency graph with lower rank, where the rank of a position is the maximal number of special edges on an incoming path. One can then inductively bound the number of possible groundings of the body to consequently bound the number of nulls generated in the head for each position.

Theorem 10. *Let Π, τ be a MV^\pm -instance program. If $\text{ve}(\Pi)$ is weakly acyclic, then $\text{OLim}(\Pi^{\text{crisp}}, D_\tau)$ is finite and of polynomial size in data complexity.*