



Seminaïve Materialisation in DatalogMTL

Dingmin Wang^(✉), Przemysław Andrzej Wałęga, and Bernardo Cuenca Grau

Department of Computer Science, University of Oxford, Oxford, UK
{dingmin.wang, przemyslaw.walega, bernardo.cuenca.grau}@cs.ox.ac.uk

Abstract. DatalogMTL is an extension of Datalog with metric temporal operators that has found applications in temporal ontology-based data access and query answering, as well as in stream reasoning. Practical algorithms for DatalogMTL are reliant on materialisation-based reasoning, where temporal facts are derived in a forward chaining manner in successive rounds of rule applications. Current materialisation-based procedures are, however, based on a naïve evaluation strategy, where the main source of inefficiency stems from redundant computations. In this paper, we propose a materialisation-based procedure which, analogously to the classical seminaïve algorithm in Datalog, aims at minimising redundant computation by ensuring that each temporal rule instance is considered at most once during the execution of the algorithm. Our experiments show that our optimised seminaïve strategy for DatalogMTL is able to significantly reduce materialisation times.

Keywords: DatalogMTL · Temporal reasoning · Materialisation

1 Introduction

DatalogMTL is a temporal rule-based language that has found a growing number of applications in ontology-based data access [9–11] and stream reasoning [20], amongst others [13, 16]. DatalogMTL extends Datalog [1, 6] with operators from metric temporal logic [12] interpreted over the rational timeline. For example, the following rule states that travellers can enter the US if they had a negative test sometime in the last 2 days ($\diamond_{[0,2]}$) and have held fully vaccinated status throughout the last 15 days ($\boxplus_{[0,15]}$):

$$\textit{Authorised}(x) \leftarrow \diamond_{[0,2]} \textit{NegativeLFT}(x) \wedge \boxplus_{[0,15]} \textit{FullyVaccinated}(x).$$

Datasets in this setting consist of temporal facts composed of a first-order fact annotated with a temporal interval, for example, $\textit{Authorised}(\textit{John})@[13, 213.5]$.

DatalogMTL is a powerful KR language and standard reasoning tasks, such as consistency and fact entailment, are PSPACE-complete in data complexity [18]. This makes efficient implementation in data-intensive applications challenging.

The most common technique of choice in scalable Datalog reasoners is materialisation (a.k.a., forward chaining) [2, 4, 5, 15]. Facts entailed by a program and

dataset are derived in successive rounds of rule applications until a fixpoint is reached; both this process and its output are often referred to as *materialisation*. The seminaïve algorithm forms the basis for efficient implementation by ensuring that each inference during rule application is only performed once, thus eliminating redundant computations. Once the materialisation has been computed, queries can be answered directly and rules are not further considered.

The use of metric temporal operators in rules, however, introduces a number of challenges for materialisation-based reasoning. First, interpretations over the rational timeline are intrinsically infinite, whereas partial materialisations computed during reasoning must be finitely represented. Second, in contrast to Datalog where materialisation naturally terminates, in DatalogMTL a fixpoint may only be reachable after infinitely many rounds of rule applications. As a matter of fact, reasoning techniques initially proposed for DatalogMTL were not materialisation-based. In particular, optimal decision procedures are automata-based [18], and reasoning is also feasible by reduction to satisfiability checking in linear temporal logic (LTL) [3]; finally, the Ontop system implements a query rewriting approach which is applicable only to non-recursive programs [9].

In our recent work [22], we proposed a materialisation-based procedure optimised for efficient application of DatalogMTL rules by means of suitable temporal indices, and where partial materialisations are succinctly represented as sets of temporal facts. We also identified a fragment of DatalogMTL [21] for which our materialisation-based procedure is guaranteed to terminate; this fragment imposes suitable restrictions which effectively disallow programs expressing ‘recursion through time’. To ensure termination in the general case for consistency and fact entailment tasks, we proposed and implemented in the MeTeoR system [22] an algorithm combining materialisation with the construction of Büchi automata, so that the use of automata-based techniques is minimised in favour of materialisation; thus, the scalability of this approach in most practical cases is critically dependent on that of its materialisation component. The materialisation-based procedure in MeTeoR is, however, based on a *naïve* strategy where the main source of inefficiency stems from redundant computations.

In this paper, we propose a seminaïve materialisation-based procedure for DatalogMTL, which can be seamlessly applied in isolation to finitely materialisable fragments [21], or used in the general case in combination with automata-based techniques [22]. As in [22], our procedure iteratively performs *materialisation steps* which compute partial materialisations consisting of temporal facts; furthermore, each materialisation step performs a round of rule applications followed by a coalescing phase where temporal facts differing only in their (overlapping) intervals are merged together. However, in contrast to [22] and analogously to the classical seminaïve algorithm for Datalog [1], our procedure aims at minimising redundant computation by considering only rule instances that involve information newly derived in the previous materialisation step. Lifting the seminaïve strategy to DatalogMTL involves significant technical challenges. In particular, rule bodies now involve metric atoms, and derived temporal facts can be coalesced with existing facts in the previous partial materialisation. As a result,

keeping track of new information and identifying the relevant rule instances to consider in each materialisation step becomes much more involved than in Datalog. We show that these difficulties can be overcome in an elegant and effective way, and propose additional optimisations aimed at further reducing redundancy for programs satisfying certain syntactic restrictions.

We have implemented our approach as an extension of MeTeoR and evaluated its performance. Our experiments show that our seminaïve strategy and the additional optimisations lead to significant reductions in materialisation times. A technical appendix containing full proofs of our technical results and the code for our implementation are available online.¹

2 Preliminaries

We recapitulate the definition of DatalogMTL interpreted under the standard continuous semantics for the rational timeline [3].

A *relational atom* is a function-free first-order atom of the form $P(\mathbf{s})$, with P a predicate and \mathbf{s} a tuple of terms. A *metric atom* is an expression given by the following grammar, where $P(\mathbf{s})$ is a relational atom, and $\diamond, \heartsuit, \boxminus, \boxplus, \mathcal{S}, \mathcal{U}$ are MTL operators indexed with intervals ϱ containing only non-negative rationals:

$$M := \top \mid \perp \mid P(\mathbf{s}) \mid \diamond_{\varrho} M \mid \heartsuit_{\varrho} M \mid \boxminus_{\varrho} M \mid \boxplus_{\varrho} M \mid M\mathcal{S}_{\varrho}M \mid M\mathcal{U}_{\varrho}M.$$

We call \diamond, \boxminus , and \mathcal{S} *past operators* and we call \heartsuit, \boxplus , and \mathcal{U} *future operators*. A *rule* is an expression of the form

$$M' \leftarrow M_1 \wedge \cdots \wedge M_n, \quad \text{for } n \geq 1, \quad (1)$$

with each M_i a metric atom, and M' is generated by the following grammar:²

$$M' ::= \top \mid P(\mathbf{s}) \mid \boxminus_{\varrho} M' \mid \boxplus_{\varrho} M'.$$

The conjunction $M_1 \wedge \cdots \wedge M_n$ in Expression (1) is the rule's *body* and M' is the rule's *head*. A rule is *forward-propagating* if it does not mention \top or \perp , mentions only past operators in the body, and only future operators in the head. A rule is *backwards-propagating*, if it satisfies analogous conditions but with past operators replaced with future operators and vice versa. A rule is *safe* if each variable in its head also occurs in the body, and this occurrence is not in a left operand of \mathcal{S} or \mathcal{U} . A *program* is a finite set of safe rules; it is forward- or backward-propagating if so are all its rules.

An expression is *ground* if it mentions no variables. A *fact* is an expression $M@_{\varrho}$ with M a ground relational atom and ϱ an interval; a *dataset* is a finite set of facts. The *coalescing* of facts $M@_{\varrho_1}$ and $M@_{\varrho_2}$, where ϱ_1 and ϱ_2 are adjacent or have a non-empty intersection, is the fact $M@_{\varrho_3}$ with ϱ_3 the union of ϱ_1 and ϱ_2 . The *grounding* $\text{ground}(\Pi, \mathcal{D})$ of program Π with respect to dataset \mathcal{D} is the set of ground rules obtained by assigning constants in Π or \mathcal{D} to variables in Π .

¹ <https://github.com/wdimmy/MeTeoR/tree/main/experiments/RR2022>.

² For presentation convenience, we disallow \perp in rule heads, which ensures satisfiability and allows us to focus on the materialisation process itself.

Table 1. Semantics of ground metric atoms

$\mathfrak{I}, t \models \top$		for each t
$\mathfrak{I}, t \models \perp$		for no t
$\mathfrak{I}, t \models \diamond_{\varrho} M$	iff	$\mathfrak{I}, t' \models M$ for some t' with $t - t' \in \varrho$
$\mathfrak{I}, t \models \oplus_{\varrho} M$	iff	$\mathfrak{I}, t' \models M$ for some t' with $t' - t \in \varrho$
$\mathfrak{I}, t \models \boxplus_{\varrho} M$	iff	$\mathfrak{I}, t' \models M$ for all t' with $t - t' \in \varrho$
$\mathfrak{I}, t \models \boxminus_{\varrho} M$	iff	$\mathfrak{I}, t' \models M$ for all t' with $t' - t \in \varrho$
$\mathfrak{I}, t \models M_1 \mathcal{S}_{\varrho} M_2$	iff	$\mathfrak{I}, t' \models M_2$ for some t' with $t - t' \in \varrho$ and $\mathfrak{I}, t'' \models M_1$ for all $t'' \in (t', t)$
$\mathfrak{I}, t \models M_1 \mathcal{U}_{\varrho} M_2$	iff	$\mathfrak{I}, t' \models M_2$ for some t' with $t' - t \in \varrho$ and $\mathfrak{I}, t'' \models M_1$ for all $t'' \in (t, t')$

An *interpretation* \mathfrak{I} specifies, for each ground relational atom M and each time point $t \in \mathbb{Q}$, whether M holds at t , in which case we write $\mathfrak{I}, t \models M$. This extends to atoms with metric operators as shown in Table 1. For an interpretation \mathfrak{I} and an interval ϱ , we define the *projection* $\mathfrak{I} \upharpoonright_{\varrho}$ of \mathfrak{I} over ϱ as the interpretation that coincides with \mathfrak{I} on ϱ and makes all relational atoms false outside ϱ . An interpretation \mathfrak{I} satisfies a fact $M@_{\varrho}$ if $\mathfrak{I}, t \models M$ for all $t \in \varrho$. Interpretation \mathfrak{I} satisfies a ground rule r if, whenever \mathfrak{I} satisfies each body atom of r at a time point t , then \mathfrak{I} also satisfies the head of r at t . Interpretation \mathfrak{I} satisfies a (non-ground) rule r if it satisfies each ground instance of r . Interpretation \mathfrak{I} is a *model* of a program Π if it satisfies each rule in Π , and it is a *model* of a dataset \mathcal{D} if it satisfies each fact in \mathcal{D} . Program Π and dataset \mathcal{D} are *consistent* if they have a model, and they *entail* a fact $M@_{\varrho}$ if each model of both Π and \mathcal{D} is a model of $M@_{\varrho}$. Each dataset \mathcal{D} has a unique least model $\mathfrak{I}_{\mathcal{D}}$, and we say that dataset \mathcal{D} *represents* interpretation $\mathfrak{I}_{\mathcal{D}}$.

The *immediate consequence operator* T_{Π} for a program Π is a function mapping an interpretation \mathfrak{I} to the least interpretation containing \mathfrak{I} and satisfying the following property for each ground instance r of a rule in Π : whenever \mathfrak{I} satisfies each body atom of r at time point t , then $T_{\Pi}(\mathfrak{I})$ satisfies the head of r at t . The successive application of T_{Π} to $\mathfrak{I}_{\mathcal{D}}$ defines a transfinite sequence of interpretations $T_{\Pi}^{\alpha}(\mathfrak{I}_{\mathcal{D}})$ for ordinals α as follows: (i) $T_{\Pi}^0(\mathfrak{I}_{\mathcal{D}}) = \mathfrak{I}_{\mathcal{D}}$, (ii) $T_{\Pi}^{\alpha+1}(\mathfrak{I}_{\mathcal{D}}) = T_{\Pi}(T_{\Pi}^{\alpha}(\mathfrak{I}_{\mathcal{D}}))$ for α an ordinal, and (iii) $T_{\Pi}^{\alpha}(\mathfrak{I}_{\mathcal{D}}) = \bigcup_{\beta < \alpha} T_{\Pi}^{\beta}(\mathfrak{I}_{\mathcal{D}})$ for α a limit ordinal. The *canonical interpretation* $\mathfrak{C}_{\Pi, \mathcal{D}}$ of Π and \mathcal{D} is the interpretation $T_{\Pi}^{\omega_1}(\mathfrak{I}_{\mathcal{D}})$, with ω_1 the first uncountable ordinal. If Π and \mathcal{D} have a model, the canonical interpretation $\mathfrak{C}_{\Pi, \mathcal{D}}$ is the least model of Π and \mathcal{D} [3].

3 Naïve Materialisation in DatalogMTL

In this section, we formulate the naïve materialisation procedure implicit in our theoretical results in [21] and implemented in the MeTeoR reasoner [22].

In order to illustrate the execution of the algorithm and discuss the inefficiencies involved, let us consider as a running example the dataset

$$\mathcal{D}_{\text{ex}} = \{R_1(c_1, c_2)@[0, 1], R_2(c_1, c_2)@[1, 2], R_3(c_2, c_3)@[2, 3], R_5(c_2)@[0, 1]\}$$

and the program Π_{ex} consisting of the following rules:

$$\begin{aligned} R_1(x, y) &\leftarrow \diamond_{[1,1]} R_1(x, y), & (r_1) \\ \boxplus_{[1,1]} R_5(y) &\leftarrow R_2(x, y) \wedge \boxplus_{[1,2]} R_3(y, z), & (r_2) \\ R_4(x) &\leftarrow \diamond_{[0,1]} R_5(x), & (r_3) \\ R_6(y) &\leftarrow R_5(y) \wedge \boxplus_{[0,2]} R_4(y) \wedge R_1(x, y). & (r_4) \end{aligned}$$

The naïve materialisation procedure applies a rule by first identifying the facts that can ground the rule body, and then determining the maximal intervals for which all the ground body atoms hold simultaneously. For instance, the procedure applies rule r_2 to \mathcal{D}_{ex} by first noticing that relational atoms $R_2(c_1, c_2)$ and $R_3(c_2, c_3)$ can be used to ground the rule body and then establishing that $[1, 1]$ is the maximal interval for which the metric atoms $R_2(c_1, c_2)$ and $\boxplus_{[1,2]} R_3(c_2, c_3)$ in the body of the relevant instance of r_2 are simultaneously true in \mathcal{D}_{ex} ; as a result, $\boxplus_{[1,1]} R_5(c_2)@[1, 1]$ can be derived, and so fact $R_5(c_2)@[2, 2]$ is added to the materialisation. In this way, the first round of rule application of the naïve materialisation procedure on Π_{ex} and \mathcal{D}_{ex} also derives fact $R_1(c_1, c_2)@[1, 2]$ using rule r_1 and fact $R_4(c_2)@[0, 2]$ using r_3 . The following set of facts is thus derived as a result of a single step of application of the rules in Π_{ex} to \mathcal{D}_{ex} :

$$\Pi_{\text{ex}}[\mathcal{D}_{\text{ex}}] = \{R_1(c_1, c_2)@[1, 2], R_4(c_2)@[0, 2], R_5(c_2)@[2, 2]\}.$$

The following definition formalises the notion of rule application.

Definition 1. *Let r be a rule of the form $M' \leftarrow M_1 \wedge \dots \wedge M_n$, for some $n \geq 1$, and let \mathcal{D} be a dataset. The set of instances for r and \mathcal{D} is defined as follows:*

$$\begin{aligned} \text{inst}_r[\mathcal{D}] &= \{(M_1\sigma@_{\varrho_1}, \dots, M_n\sigma@_{\varrho_n}) \mid \sigma \text{ is a substitution and, for each} \\ &i \in \{1, \dots, n\}, \varrho_i \text{ is a subset-maximal interval such that } \mathcal{D} \models M_i\sigma@_{\varrho_i}\}. \end{aligned}$$

The set $r[\mathcal{D}]$ of facts derived by r from \mathcal{D} is defined as follows:

$$\begin{aligned} r[\mathcal{D}] &= \{M\sigma@_{\varrho} \mid \sigma \text{ is a substitution, } M \text{ is the single relational atom in } M'\sigma, \\ &\text{and there exists } (M_1\sigma@_{\varrho_1}, \dots, M_n\sigma@_{\varrho_n}) \in \text{inst}_r[\mathcal{D}] \text{ such that } \varrho \text{ is the unique} \\ &\text{subset-maximal interval satisfying } M'\sigma@_{(\varrho_1 \cap \dots \cap \varrho_n)} \models M\sigma@_{\varrho}\}. \end{aligned} \quad (2)$$

The set of facts derived from \mathcal{D} by one-step application of Π is

$$\Pi[\mathcal{D}] = \bigcup_{r \in \Pi} r[\mathcal{D}]. \quad (3)$$

Procedure 1: Naïve(Π, \mathcal{D})**Input:** A program Π and a dataset \mathcal{D} **Output:** A dataset representing the canonical interpretation $\mathfrak{C}_{\Pi, \mathcal{D}}$

```

1 Initialise  $\mathcal{N}$  to  $\emptyset$  and  $\mathcal{D}'$  to  $\mathcal{D}$ ;
2 loop
3    $\mathcal{N} := \Pi(\mathcal{D}')$ ;                                     // derive new facts
4    $\mathcal{C} := \mathcal{D}' \uplus \mathcal{N}$ ;                               // coalesce with the new facts
5   if  $\mathcal{C} = \mathcal{D}'$  then return  $\mathcal{D}'$ ;
6    $\mathcal{D}' := \mathcal{C}$ ;

```

Once rule application has been completed and facts $\Pi_{\text{ex}}[\mathcal{D}_{\text{ex}}]$ have been derived, the partial materialisation $\mathcal{D}_{\text{ex}}^1$ that will be passed on to the next materialisation step is obtained as $\mathcal{D}_{\text{ex}}^1 = \mathcal{D}_{\text{ex}} \uplus \Pi_{\text{ex}}[\mathcal{D}_{\text{ex}}]$ by coalescing facts in \mathcal{D}_{ex} and $\Pi_{\text{ex}}[\mathcal{D}_{\text{ex}}]$, where the coalescing operator \uplus is semantically defined next.

Definition 2. For datasets \mathcal{D}_1 and \mathcal{D}_2 , we define $\mathcal{D}_1 \uplus \mathcal{D}_2$ as the dataset consisting of all relational facts $M @ \varrho$ such that $\mathcal{D}_1 \cup \mathcal{D}_2 \models M @ \varrho$ and $\mathcal{D}_1 \cup \mathcal{D}_2 \not\models M @ \varrho'$, for each ϱ' with $\varrho \subsetneq \varrho'$.

The use of coalescing makes sure that intervals associated to facts are maximal. In our example, facts $R_1(c_1, c_2) @ [0, 1]$ in \mathcal{D}_{ex} and $R_1(c_1, c_2) @ [1, 2]$ are coalesced, so we have $R_1(c_1, c_2) @ [0, 2]$ in $\mathcal{D}_{\text{ex}} \uplus \Pi_{\text{ex}}[\mathcal{D}_{\text{ex}}]$. Thus,

$$\mathcal{D}_{\text{ex}}^1 = \{R_1(c_1, c_2) @ [0, 2], R_2(c_1, c_2) @ [1, 2], R_3(c_2, c_3) @ [2, 3], \\ R_5(c_2) @ [0, 1], R_4(c_2) @ [0, 2], R_5(c_2) @ [2, 2]\}.$$

In the second round, rules are applied to $\mathcal{D}_{\text{ex}}^1$. The application of r_1 derives fact $R_1(c_1, c_2) @ [1, 3]$ (from $R_1(c_1, c_2) @ [0, 2]$) and the application of r_2 rederives a redundant fact $R_5(c_2) @ [2, 2]$. In contrast to the previous step, rule r_4 can now be applied to derive the new fact $R_6(c_2) @ [2, 2]$. Finally, the application of r_3 derives the new fact $R_4(c_2) @ [2, 3]$ and rederives the redundant fact $R_4(c_2) @ [0, 2]$.

The procedure then coalesces fact $R_1(c_1, c_2) @ [1, 3]$ with $R_1(c_1, c_2) @ [0, 2]$ to obtain $R_1(c_1, c_2) @ [0, 3]$; similarly, $R_4(c_2) @ [2, 3]$ is coalesced with $R_4(c_2) @ [0, 2]$ to obtain $R_4(c_2) @ [0, 3]$. Thus, the second step yields the following partial materialisation:

$$\mathcal{D}_{\text{ex}}^2 = (\mathcal{D}_{\text{ex}}^1 \setminus \{R_4(c_2) @ [0, 2], R_1(c_1, c_2) @ [0, 2]\}) \cup \\ \{R_1(c_1, c_2) @ [0, 3], R_6(c_2) @ [2, 2], R_4(c_2) @ [0, 3]\}.$$

In the third materialisation step, rules are applied to $\mathcal{D}_{\text{ex}}^2$, and derive the new fact $R_1(c_1, c_2) @ [1, 4]$, as well as redundant facts such as $R_5(c_2) @ [2, 2]$, $R_4(c_2) @ [0, 2]$, $R_4(c_2) @ [2, 3]$, and $R_6(c_2) @ [2, 2]$. The procedure will then continue completing subsequent materialisation steps and stopping only if a fixpoint is reached.

Procedure 1 formalises this naïve materialisation strategy. As discussed, each iteration of the main loop captures a single materialisation step consisting of a

round of rule application (c.f. Line 3) followed by the coalescing of relevant facts (c.f. Line 4). The resulting partial materialisation passed on to the following materialisation stem is stored as a dataset \mathcal{D}' (c.f. Line 6). The procedure stops when a materialisation step does not derive any new facts, in which case a fixpoint has been reached (c.f. Line 5). In our example, materialisation will continue to recursively propagate the relational fact $R_1(c_1, c_2)$ throughout the infinite timeline, and the procedure will not terminate as a result. Furthermore, the number of redundant computations will increase in each subsequent materialisation step.

It is worth recalling that, even in cases where materialisation does not reach a fixpoint, it still constitutes a key component of terminating algorithms such as that implemented in MeTeoR [22]. Therefore, the performance challenges stemming from redundant computations remain a very significant issue in practice.

4 Seminaïve Evaluation

Seminaïve rule evaluation is the technique of choice for eliminating redundant computations in Datalog-based systems. The main idea is to keep track of newly derived facts in each materialisation step by storing them in a set Δ , and to make sure that rule applications in the following materialisation step involve at least one fact in Δ . In this way, the procedure considers each rule instance *at most once* throughout its entire execution and it is said to enjoy the *non-repetition* property. Note, however, that the same fact can still be derived multiple times by *different* rule instances; this type of redundancy is difficult to prevent and is not addressed by the standard seminaïve strategy.

Our aim in this section is to lift seminaïve rule evaluation to the setting of DatalogMTL. As discussed in Sect. 3 on our running example, a rule instance can be considered multiple times in our setting; for example, the instance $(R_2(c_1, c_2)@[1, 2], \boxplus_{[1,2]}R_3(c_2, c_3)@[1, 1])$ of r_2 is considered in both the first and second materialisation steps to derive $R_5(c_2)@[2, 2]$ twice since the naïve procedure cannot detect that facts $R_2(c_1, c_2)@[1, 2]$ and $R_3(c_2, c_3)@[2, 3]$ used to instantiate r_2 in the second step had previously been used to instantiate r_2 . Preventing such redundant computations, however, involves certain challenges. First, by including in Δ just the newly derived facts as in Datalog, we could overlook relevant information obtained by coalescing newly derived facts with previously derived ones. Second, restricting application to relevant rule instances requires taking into account the semantics of metric operators in rule bodies.

Procedure 2 extends the seminaïve strategy to the setting of DatalogMTL while overcoming the aforementioned difficulties. Analogously to the naïve approach, each iteration of the main loop captures a single materialisation step consisting of a round of rule applications followed by the coalescing of relevant facts; as before, dataset \mathcal{D}' stores the partial materialisation resulting from each iteration and is initialised as the input dataset, whereas dataset \mathcal{N} stores the facts obtained as a result of rule application and is initialised as empty.

Following the rationale behind the seminaïve strategy for Datalog, newly derived information in each materialisation step is now stored as a dataset Δ ,

Procedure 2: Seminaïve(Π, \mathcal{D})

Input: A program Π and a dataset \mathcal{D}

Output: A dataset representing the canonical interpretation $\mathfrak{C}_{\Pi, \mathcal{D}}$

```

1 Initialise  $\mathcal{N}$  to  $\emptyset$ , and both  $\Delta$  and  $\mathcal{D}'$  to  $\mathcal{D}$ ;
2 loop
3    $\mathcal{N} := \Pi[\mathcal{D}' : \Delta]$ ;
4    $\mathcal{C} := \mathcal{D}' \uplus \mathcal{N}$ ;
5    $\Delta := \{M@_q \in \mathcal{C} \mid M@_q \text{ entails some fact in } \mathcal{N} \parallel \mathcal{D}'\}$ ;
6   if  $\Delta = \emptyset$  then return  $\mathcal{D}'$ ;
7    $\mathcal{D}' := \mathcal{C}$ ;
```

which is initialised as the input dataset \mathcal{D} and which is suitably maintained in each iteration; furthermore, Procedure 2 ensures in Line 3 that only rule instances, for which it is essential to involve facts from Δ (as formalised in the following definition) are taken into account during rule application.

Definition 3. Let r be a rule of the form $M' \leftarrow M_1 \wedge \dots \wedge M_n$, for some $n \geq 1$, and let \mathcal{D} and Δ be datasets. The set of instances for r and \mathcal{D} relative to Δ is defined as follows:

$$\text{inst}_r[\mathcal{D} : \Delta] = \{(M_1 \sigma@_{q_1}, \dots, M_n \sigma@_{q_n}) \in \text{inst}_r[\mathcal{D}] \mid \mathcal{D} \setminus \Delta \not\models M_i \sigma@_{q_i}, \text{ for some } i \in \{1, \dots, n\}\}. \quad (4)$$

The set $r[\mathcal{D} : \Delta]$ of facts derived by r from \mathcal{D} relative to Δ is defined analogously to $r[\mathcal{D}]$ in Definition 1, with the exception that $\text{inst}_r[\mathcal{D}]$ is replaced with $\text{inst}_r[\mathcal{D} : \Delta]$ in Expression (2). Finally, the set $\Pi[\mathcal{D} : \Delta]$ of facts derived from \mathcal{D} by one-step seminaïve application of Π is defined as $\Pi[\mathcal{D}]$ in Expression (3), by replacing $r[\mathcal{D}]$ with $r[\mathcal{D} : \Delta]$.

In each materialisation step, Procedure 2 exploits Definition 3 to identify as relevant the subset of rule instances where some conjunct is ‘new’, in the sense that it cannot be entailed without the facts in Δ . The facts derived by such relevant rule instances in each iteration are stored in set \mathcal{N} (c.f. Line 3).

As in the naïve approach, rule application is followed by a coalescing step where the partial materialisation is updated with the facts derived from rule application (c.f. Line 4). In contrast to the naïve approach, however, Procedure 2 needs to maintain set Δ to ensure that it captures only new facts. This is achieved in Line 5, where a fact in the updated partial materialisation is considered new if it entails a fact in \mathcal{N} that was not already entailed by the previous partial materialisation. This is formalised with the following notion of ‘semantic’ difference between temporal datasets.

Definition 4. Let \mathcal{D}_1 and \mathcal{D}_2 be datasets. We define $\mathcal{D}_1 \parallel \mathcal{D}_2$ as the dataset consisting of all relational facts $M@_q$ such that $M@_q \in \mathcal{D}_1$ and $\mathcal{D}_2 \not\models M@_q$.

The procedure terminates in Line 6 if Δ is empty. Otherwise, the procedure carries over the updated partial materialisation and the set of newly derived facts to the next materialisation step.

We next illustrate the execution of the procedure on \mathcal{D}_{ex} and Π_{ex} . In the first materialisation step, all input facts are considered as newly derived (i.e., $\Delta = \mathcal{D}$) and hence $\mathcal{N} = \Pi[\mathcal{D}' : \Delta] = \Pi(\mathcal{D}')$ and the result of coalescing coincides with the partial materialisation computed by the naïve procedure (i.e., $\mathcal{C} = \mathcal{D}_{\text{ex}}^1$). Then, the procedure identifies as new all facts in \mathcal{N} (i.e., $\Delta = \mathcal{N}$). In the second step, rule evaluation in Line 3 no longer considers the redundant instance of r_2 consisting of fact $R_2(c_1, c_2)@[1, 2]$ and metric atom $\boxplus_{[1,2]}R_3(c_2, c_3)@[1, 2]$ since they are respectively entailed by facts $R_2(c_1, c_2)[1, 2]$ and $R_3(c_2, c_3)[2, 3]$ in $\mathcal{D}' \setminus \Delta$. Finally, the procedure also disregards the redundant instance of r_3 re-deriving fact $R_4(c_2)[0, 2]$. In contrast, all non-redundant facts derived by the naïve strategy are also derived by the seminaïve procedure and after coalescing dataset $\mathcal{C} = \mathcal{D}_{\text{ex}}^2$. Set Δ is now updated as follows:

$$\Delta = \{R_1(c_2, c_2)@[0, 3], R_6(c_2)@[2, 2], R_4(c_2)@[0, 3]\}.$$

In particular, note that Δ contains the coalesced fact $R_4(c_2)@[0, 3]$ rather than fact $R_4(c_2)@[2, 3]$ derived from rule application. Datasets Δ and $\mathcal{D}' = \mathcal{D}_{\text{ex}}^2$ are passed on to the third materialisation step, where all redundant computations identified in Sect. 3 are avoided with the only exception of fact $R_6(c_2)@[2, 2]$, which is re-derived using the instance of r_4 consisting of facts $R_5(c_2)@[2, 2]$, $R_1(c_2, c_2)@[0, 3]$ and metric atom $\boxminus_{[0,2]}R_4(c_2)@[2, 3]$. Note that this is a new instance which was not used in previous iterations, and hence the non-repetition property remains true. Note also that, as with the naïve strategy, our seminaïve procedure does not terminate on our running example.

We conclude this section by establishing correctness of our procedure. To this end we next show that, upon completion of the k -th iteration of the main loop (for any k), the partial materialisation \mathcal{D}' passed on to the next iteration represents the interpretation $T_H^k(\mathcal{J}_{\mathcal{D}})$ obtained by applying k times the immediate consequence operator T_H for the input program Π to the interpretation $\mathcal{J}_{\mathcal{D}}$ representing the input dataset \mathcal{D} . This provides a precise correspondence between the procedure's syntactic operations and the semantics of fixpoint computation.

Soundness relies on the observation that rule instances processed by seminaïve evaluation are also processed by the naïve evaluation; thus, $\text{inst}_r[\mathcal{D}' : \Delta] \subseteq \text{inst}_r[\mathcal{D}]$, for each r , \mathcal{D} , and Δ . As a result, each fact derived by the seminaïve evaluation is also derived by the naïve evaluation.

Theorem 1 (Soundness). *Consider Procedure 2 running on input Π and \mathcal{D} . Upon the completion of the k th (for some $k \in \mathbb{N}$) iteration of the loop of Procedure 2, it holds that $\mathcal{J}_{\mathcal{D}'} \subseteq T_H^k(\mathcal{J}_{\mathcal{D}})$.*

Completeness is proved by induction on the number k of iterations of the main loop. In particular, we show that if $T_H^k(\mathcal{J}_{\mathcal{D}})$ satisfies a new fact $M@t$, then there must be a rule r and an instance in $\text{inst}_r[\mathcal{D}' : \Delta]$ witnessing the derivation of $M@t$; otherwise, the fact would hold already in $T_H^{k-1}(\mathcal{J}_{\mathcal{D}})$. Hence, each fact satisfied by $T_H^k(\mathcal{J}_{\mathcal{D}})$ is derived in the k th iteration of our procedure.

Theorem 2 (Completeness). *Consider Procedure 2 running on input Π and \mathcal{D} . For each $k \in \mathbb{N}$, upon the completion of the k th iteration of the loop of Procedure 2, it holds that $T_{\Pi}^k(\mathcal{I}_{\mathcal{D}}) \subseteq \mathcal{I}_{\mathcal{D}'}$.*

5 Optimised Seminaïve Evaluation

Although the seminaïve procedure enjoys the non-repetition property, it can still re-derive facts that were already obtained in previous materialisation steps, thus incurring in a potentially large number of redundant computations. In particular, as discussed in Sect. 4, fact $R_6(c_2)@[2, 2]$ is re-derived using rule r_4 in the third materialisation step of our running example, and it will also be re-derived in all subsequent materialisation steps (by different instances of rule r_4).

In this section, we present an optimised variant of our seminaïve procedure which further reduces the number of redundant computations performed during materialisation. The main idea is to disregard rules during the execution of the procedure as soon as we can be certain that their application will never derive new facts in subsequent materialisation steps. In our example, rule r_4 can be discarded after the second materialisation step as its application will only continue to re-derive fact $R_6(c_2)@[2, 2]$ in each materialisation step.

To this end, we will exploit the distinction between recursive and non-recursive predicates in a program, as defined next.

Definition 5. *The dependency graph of program Π is the directed graph with a vertex v_P for each predicate P in Π and an edge (v_Q, v_R) whenever there is a rule in Π mentioning Q in the body and R in the head. Predicate P is recursive (in Π) if the dependency graph has a path containing a cycle and ending in v_P ; otherwise P is non-recursive. A metric atom is non-recursive in Π if so are all its predicates; otherwise it is recursive. The (non-)recursive fragment of Π is the subset of rules in Π with (non-)recursive atoms in heads.*

In contrast to recursive predicates, for which new facts can be derived in each materialisation step, the materialisation of non-recursive predicates will be completed after linearly many materialisation steps; from then on, the rules will no longer derive any new facts involving these predicates.

This observation can be exploited to optimise seminaïve evaluation. Assume that the procedure has fully materialised all non-recursive predicates in the input program. At this point, we can safely discard all non-recursive rules; furthermore, we can also discard a recursive rule r with a non-recursive body atom M if the current partial materialisation does not entail any grounding of M (in this case, r cannot apply in further materialisation steps). An additional optimisation applies to forward-propagating programs, where rules cannot propagate information ‘backwards’ along the timeline; in this case, we can compute the maximal time points for which each non-recursive body atom in r may possibly hold, select the minimum t_r amongst such values, and discard r as soon as we can determine that the materialisation up to time point t_r has been fully completed.

Procedure 3: OptimisedSeminaïve(Π, \mathcal{D})

Input: A program Π and a dataset \mathcal{D}
Output: A dataset representing the canonical interpretation $\mathfrak{C}_{\Pi, \mathcal{D}}$

```

1 Initialise  $\mathcal{N}$  to  $\emptyset$ ,  $\Delta$  and  $\mathcal{D}'$  to  $\mathcal{D}$ ,  $\Pi'$  to  $\Pi$ , flag to 0, and  $S_r$  (for each  $r \in \Pi$ )
   to the set of body atoms in  $r$  that are non-recursive in  $\Pi$ ;
2 loop
3    $\mathcal{N} := \Pi'[\mathcal{D}' : \Delta]$ ;
4    $\mathcal{C} := \mathcal{D}' \uplus \mathcal{N}$ ;
5    $\Delta := \{M@q \in \mathcal{C} \mid M@q \text{ entails some fact in } \mathcal{N} \parallel \mathcal{D}'\}$ ;
6   if  $\Delta = \emptyset$  then return  $\mathcal{D}'$ ;
7   if flag = 0 &  $\mathcal{D}'$  and  $\mathcal{C}$  entail the same facts with non-recursive predicates
   then
8     Set flag to 1 and  $\Pi'$  to the recursive fragment of  $\Pi$ ;
9     for each  $r \in \Pi'$  do
10      if there is  $M \in S_r$  such that  $\mathcal{D}' \not\models M\sigma@t$ , for each substitution  $\sigma$ 
        and time point  $t$  then  $\Pi' := \Pi' \setminus \{r\}$ ;
11  if flag = 1 and  $\Pi'$  is forward propagating then
12    for each  $r \in \Pi'$  and each  $M \in S_r$  do
13       $t_{max}^M :=$  maximum right endpoint amongst all intervals  $q$  satisfying
         $\mathcal{D}' \models M\sigma@q$ , for some substitution  $\sigma$ ;
14       $t_r :=$  minimum value in  $\{t_{max}^M \mid M \in S_r\}$ ;
15      if  $\mathcal{D}'$  and  $\mathcal{C}$  coincide on all facts over intervals  $q$  satisfying  $q^+ \leq t_r$ 
        then  $\Pi' := \Pi' \setminus \{r\}$ ;
16   $\mathcal{D}' := \mathcal{C}$ ;

```

The materialisation of the non-recursive predicates R_2 , R_3 , R_4 , and R_5 of our running example is complete after two materialisation steps. Hence, at this point we can disregard rules r_2 and r_3 and focus on the recursive forward-propagating rules r_1 and r_4 . Furthermore, the maximum time point at which R_4 and R_5 can hold is 3 and 2, respectively, and hence $t_{r_4} = 2$; thus, upon completion of the second materialisation step we can be certain that R_1 has been materialised up to t_r and we can also discard r_6 . In subsequent materialisation steps we can apply only rule r_1 , thus avoiding many redundant computations.

Procedure 3 implements these ideas by extending seminaïve materialisation. In each materialisation step, the procedure checks whether all non-recursive predicates have been fully materialised (c.f. Line 7), in which case it removes all non-recursive rules in the input program as well as all recursive rules with an unsatisfied non-recursive body atom (c.f. Lines 7–10). It also sets a flag to 1, which activates the additional optimisation for forward propagating programs, which is applied in Lines 11–15 whenever possible.

We conclude this section by establishing correctness of our procedure. We first observe that, as soon as the algorithm switches the flag to 1, we can be certain that all non-recursive predicates have been fully materialised.

Lemma 1. *Consider Procedure 3 running on input Π and \mathcal{D} and let Π_{nr} be the non-recursive fragment of Π . If $flag = 1$, then $\mathfrak{C}_{\Pi_{nr}, \mathcal{D}} \subseteq \mathfrak{I}_{\mathcal{D}'}$.*

Next, we show how we can detect if a forward-propagating program has completed materialisation of all facts (also with recursive predicates) up to a given time point. Namely it suffices to check that two consecutive partial materialisations satisfy the same facts up to a given time point. Note that our procedure checks this condition syntactically in Line 15.

Lemma 2. *If $\mathfrak{I}_{\mathcal{D}} \upharpoonright_{(-\infty, t]} = T_{\Pi}(\mathfrak{I}_{\mathcal{D}}) \upharpoonright_{(-\infty, t]}$, for a forward propagating program Π , dataset \mathcal{D} , and time point t , then $\mathfrak{I}_{\mathcal{D}} \upharpoonright_{(-\infty, t]} = \mathfrak{C}_{\Pi, \mathcal{D}} \upharpoonright_{(-\infty, t]}$.*

We can use this lemma to show that each rule discarded in Lines 10 and 15 can be safely ignored as it will have no effect in subsequent materialisation steps.

Lemma 3. *If in Procedure 3 a rule r is removed from Π' in Line 10 or in Line 15, then $\mathfrak{C}_{\Pi', \mathcal{D}' \uplus \mathcal{N}} = \mathfrak{C}_{\Pi' \setminus \{r\}, \mathcal{D}' \uplus \mathcal{N}}$.*

Finally, using Lemmas 1 and 3 together with the soundness and completeness of our seminaïve evaluation (established in Theorems 2 and 1), we can show soundness and completeness of the optimised version of the procedure.

Theorem 3 (Soundness and Completeness). *Consider Procedure 3 running on input Π and \mathcal{D} . For each $k \in \mathbb{N}$, the partial materialisation \mathcal{D}' obtained upon completion of the k th iteration of the main loop represents the interpretation $T_{\Pi}^k(\mathfrak{I}_{\mathcal{D}})$.*

We conclude by observing that our optimisation for forward-propagating programs in Lines 11–15 can be modified in a straightforward way to account also for backwards-propagating programs, as these two cases are symmetric.

6 Evaluation

We have implemented Procedures 2 and 3 as an extension of our open-source MeTeoR reasoner [22], which so-far implemented only the naïve strategy from Procedure 1 to perform materialisation.

For evaluation, we have considered the temporal extension of the Lehigh University Benchmark (LUBM) [8] used in previous evaluations of MeTeoR [22]. The benchmark provides a DatalogMTL program consisting of 56 Datalog rules obtained from the OWL 2 RL fragment of LUBM’s ontology plus 29 temporal rules involving recursion and covering all metric operators of DatalogMTL. To make materialisation more challenging, we have included additional body atoms in some of the temporal rules. The benchmark also provides an extension of LUBM’s data generator which randomly assigns intervals non-temporal facts. We used nine datasets \mathcal{D}_1 – \mathcal{D}_9 , each consisting of 10 million facts, but with an increasing number of constants occurring in these facts (and thus with a smaller number of intervals per relational fact), namely, these datasets contain 0.8, 1.0,

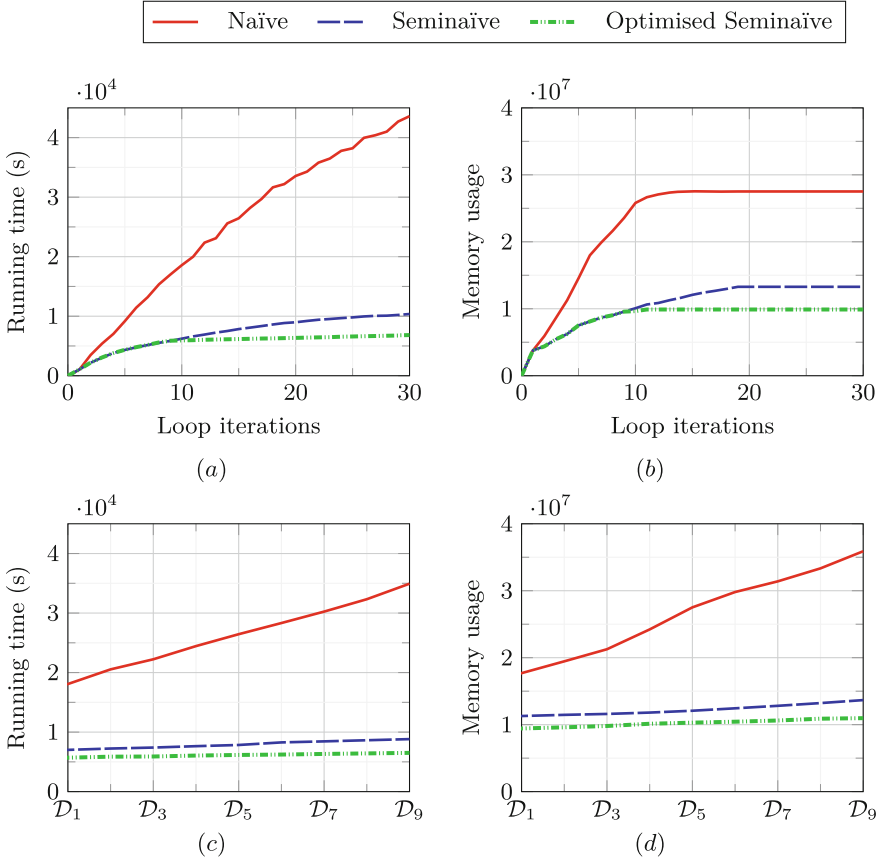


Fig. 1. Experimental results for dataset \mathcal{D}_5 in sub-figures (a) and (b), and for the first 15 iterations for datasets \mathcal{D}_1 – \mathcal{D}_9 in sub-figures (c) and (d)

1.2, 1.3, 2.1, 2.5, 5.2, 10.1, and 15.8 million constants, respectively. We compared running time and memory requirements (maximal number of stored facts) of our procedures with that of the naïve approach as depicted in Fig. 1. Experiments were conducted on a Dell PowerEdge R730 server with 512 GB RAM and two Intel Xeon E5-2640 2.6 GHz processors running Fedora 33, kernel version 5.8.17.

Figures 1 (a) and (b) show time and memory usage on a single dataset \mathcal{D}_5 through the first 30 iterations of the procedures. As we can see, the seminaïve procedure significantly outperforms the naïve approach both in terms of running time and memory consumption, especially as materialisation progresses. In turn, the optimised seminaïve approach is able to start disregarding rules after 9 materialisation steps, and at that point it starts outperforming the basic seminaïve procedure. We can also observe that, at this point, many of the predicates have been materialised already and the number of new facts generated in each further step is very small (thus, memory consumption stops growing). Despite

this, the naïve algorithm continues to perform a large number of redundant computations, and hence its running time increases at a similar rate as before; in contrast, the optimised seminaïve approach avoids most of this redundancy and subsequent materialisation steps are completed very efficiently.

Figures 1 (c) and (d) summarise our results for 15 materialisation steps on datasets \mathcal{D}_1 – \mathcal{D}_9 with increasing numbers of constants. We can observe that both the time and memory consumption in the naïve approach increase linearly (and in a significant way) with the number of constants. Indeed, by increasing the number of constants, we are also increasing the number of ground rule instances to be examined. In contrast, the effect on both of our seminaïve procedures is much less noticeable as they can quickly disregard irrelevant instances.

7 Conclusion and Future Work

In this paper, we have presented an optimised seminaïve materialisation procedure for DatalogMTL, which can efficiently materialise complex recursive programs and large datasets involving millions of temporal facts.

We see many exciting avenues for future research. First, DatalogMTL has been extended with stratified negation-as-failure [7] and our seminaïve procedure could be extended accordingly. It would also be interesting to consider seminaïve evaluation for reasoning under alternative semantics for DatalogMTL such as the integer semantics [19] or the point-wise semantics [17]. We are also working on blocking conditions that exploit the periodic structure of canonical models to ensure termination of materialisation-based reasoning. Finally, incremental materialisation-based reasoning has been studied in context of Datalog [14], and it would be interesting to lift such approaches to the DatalogMTL setting.

Acknowledgments. This work was supported by the EPSRC project OASIS (EP/S032347/1), the EPSRC project UK FIRES (EP/S019111/1), and the SIRIUS Centre for Scalable Data Access, and Samsung Research UK.

References

1. Abiteboul, S., Hull, R., Vianu, V.: *Foundations of Databases*, vol. 8. Addison-Wesley, Reading (1995)
2. Bellomarini, L., Sallinger, E., Gottlob, G.: The vatalog system: Datalog-based reasoning for knowledge graphs. *Proc. VLDB Endow.* **11**(9), 975–987 (2018)
3. Brandt, S., Kalaycı, E.G., Ryzhikov, V., Xiao, G., Zakharyashev, M.: Querying log data with metric temporal logic. *J. Artif. Intell. Res.* **62**, 829–877 (2018)
4. Bry, F., et al.: Foundations of rule-based query answering. In: *Reasoning Web*, pp. 1–153 (2007)
5. Carral, D., Dragoste, I., González, L., Jacobs, C.J.H., Krötzsch, M., Urbani, J.: Vlog: a rule engine for knowledge graphs. In: *Proceedings of ISWC*, pp. 19–35 (2019)
6. Ceri, S., Gottlob, G., Tanca, L.: What you always wanted to know about Datalog (and never dared to ask). *IEEE TKDE* **1**(1), 146–166 (1989)

7. Cucala, D.J.T., Wałęga, P.A., Cuenca Grau, B., Kostylev, E.V.: Stratified negation in Datalog with metric temporal operators. In: Proceedings of AAAI, pp. 6488–6495 (2021)
8. Guo, Y., Pan, Z., Heflin, J.: LUBM: a benchmark for OWL knowledge base systems. *J. Web Semant.* **3**(2–3), 158–182 (2005)
9. Kalaycı, E.G., Xiao, G., Ryzhikov, V., Kalaycı, T.E., Calvanese, D.: Onto-temporal: a tool for ontology-based query answering over temporal data. In: Proceedings of CIKM, pp. 1927–1930 (2018)
10. Kikot, S., Ryzhikov, V., Wałęga, P.A., Zakharyashev, M.: On the data complexity of ontology-mediated queries with MTL operators over timed words. In: Proceedings of DL (2018)
11. Koopmann, P.: Ontology-based query answering for probabilistic temporal data. In: Proceedings of AAAI, pp. 2903–2910 (2019)
12. Koymans, R.: Specifying real-time properties with metric temporal logic. *J. R Time Syst.* **2**(4), 255–299 (1990)
13. Mori, M., Papotti, P., Bellomarini, L., Giudice, O.: Neural machine translation for fact-checking temporal claims. In: Proceedings of FEVER, p. 78 (2022)
14. Motik, B., Nenov, Y., Piro, R., Horrocks, I.: Maintenance of Datalog materialisations revisited. *Artif. Intell.* **269**, 76–136 (2019)
15. Motik, B., Nenov, Y., Piro, R., Horrocks, I., Olteanu, D.: Parallel materialisation of Datalog programs in centralised, main-memory RDF systems. In: Proceedings of AAAI (2014)
16. Nissl, M., Sallinger, E.: Modelling smart contracts with datalogmtl. In: Ramanath, M., Palpanas, T. (eds.) Proceedings of the Workshops of the EDBT/ICDT. CEUR, vol. 3135. CEUR-WS.org (2022)
17. Ryzhikov, V., Wałęga, P.A., Zakharyashev, M.: Data complexity and rewritability of ontology-mediated queries in metric temporal logic under the event-based semantics. In: Proceedings of IJCAI, pp. 1851–1857 (2019)
18. Wałęga, P.A., Cuenca Grau, B., Kaminski, M., Kostylev, E.V.: DatalogMTL: computational complexity and expressive power. In: Proceedings of IJCAI, pp. 1886–1892 (2019)
19. Wałęga, P.A., Cuenca Grau, B., Kaminski, M., Kostylev, E.V.: DatalogMTL over the integer timeline. In: Proceedings of KR, pp. 768–777 (2020)
20. Wałęga, P.A., Kaminski, M., Cuenca Grau, B.: Reasoning over streaming data in metric temporal Datalog. In: Proceedings of AAAI, pp. 3092–3099 (2019)
21. Wałęga, P.A., Zawidzki, M., Cuenca Grau, B.: Finitely materialisable Datalog programs with metric temporal operators. In: Proceedings of KR (2021)
22. Wang, D., Hu, P., Wałęga, P.A., Grau, B.C.: MeTeoR: practical reasoning in Datalog with metric temporal operators. In: Proceedings of AAAI (2022)