

Capturing Homomorphism-Closed Decidable Queries with Existential Rules

Camille Bourgaux¹, David Carral², Markus Krötzsch³, Sebastian Rudolph³, Michaël Thomazo¹

¹ DI ENS, ENS, CNRS, PSL University & Inria, Paris, France

² LIRMM, Inria, University of Montpellier, CNRS, Montpellier, France

³ Technische Universität Dresden, Dresden, Germany

{camille.bourgaux, david.carral, michael.thomazo}@inria.fr,

{markus.kroetzsch, sebastian.rudolph}@tu-dresden.de

Abstract

Existential rules are a very popular ontology-mediated query language for which the chase represents a generic computational approach for query answering. It is straightforward that existential rule queries exhibiting chase termination are decidable and can only recognize properties that are preserved under homomorphisms. In this paper, we show the converse: every decidable query that is closed under homomorphism can be expressed by an existential rule set for which the standard chase universally terminates. Membership in this fragment is not decidable, but we show via a diagonalisation argument that this is unavoidable.

1 Introduction

At the core of contemporary logic-based knowledge representation is the concept of *querying* data sources, often using elaborate query formalisms that allow for taking background knowledge into account. The classical decision problem related to such knowledge-aware querying is *Boolean query entailment*. From an abstract point of view, a Boolean query identifies a class of *databases* \mathcal{D} – those that satisfy the query, i.e., to which the query “matches”. This view allows us to define and investigate properties of (abstract) queries independently from the syntax used to specify them. Such properties can be structural (morphisms, closure properties) or computational (decidability, complexity).

A very popular querying formalism are *existential rules*, also referred to as *tuple-generating dependencies*. It is straightforward that the class of databases satisfying some existential rule query is closed under homomorphisms and recursively enumerable. Conversely, it was established that every homomorphism-closed query that is recursively enumerable can be expressed using existential rules (Rudolph and Thomazo 2015). That is, plain existential rules already realize their full potential; further syntactic extensions within these boundaries do not enhance expressivity.

For questions related to automated deduction, however, a more restricted requirement than recursive enumerability is of central interest: decidability. Therefore, the crucial question we tackle in this paper is:

Can we characterize an existential rules fragment capable of expressing every decidable homomorphism-closed query?

The generic computational paradigm for existential rules, the *chase* (Beeri and Vardi 1984), is based on repetitive, for-

ward-chaining rule application, starting from the database. As this may cause the iterated introduction of new domain elements, this procedure is not guaranteed to terminate – yet, termination is a crucial criterion for decidability. The chase comes in several variants, mainly differing in their (increasingly thorough) mechanisms to prevent unnecessary rule applications: While the *Skolem* chase (Marnette 2009) essentially just avoids duplicate rule applications, the *standard* (Fagin et al. 2005) and the *core* chase (Deutsch, Nash, and Remmel 2008) check for redundancy on a local and global level, respectively.

The class of existential rule sets with terminating¹ Skolem chase has already been weighed and found wanting: it only comprises those queries that are already expressible in plain Datalog – and hence can be evaluated in polynomial time (Marnette 2009; Krötzsch and Rudolph 2011; Zhang, Zhang, and You 2015). For the standard-chase-terminating and the core-chase-terminating existential rules classes, on the other hand, we only know that the former is contained in the latter (Grahne and Onet 2018), but little more than that (Krötzsch, Marx, and Rudolph 2019). In this paper, we clarify the situation significantly by showing the following:

Standard-chase-terminating existential rules capture the class of all decidable homomorphism-closed queries.

Notably, this implies that standard-chase-terminating and core-chase-terminating existential rule queries are equally expressive and no decidable enhancement of this formalism that preserves homomorphism-closedness (e.g. by allowing disjunction in rule heads) can be strictly more expressive.

As a downside, the existential rules fragment thus identified is not even semi-decidable, but we show via a diagonalisation argument that this downside is, in fact, unavoidable.

Additional proofs and details are given in the appendix of (Bourgaux et al. 2021).

2 Preliminaries

Rules We consider first-order formulas over countably infinite sets Vars of variables and Preds of *predicates*, where each $p \in \text{Preds}$ has an arity $Ar(p) \geq 0$. Lists of variables are denoted $\vec{x} = x_1, \dots, x_k$ and will be treated like sets when order is not relevant. An *atom* is an expression $p(\vec{x})$ with $p \in \text{Preds}$ and $|\vec{x}| = Ar(p)$.

¹We always mean universal termination, i.e., for every database.

The fragment of *disjunctive existential rules* consists of formulae of the form:

$$\forall \vec{x}. \left(\beta[\vec{x}] \rightarrow \bigvee_{i=1}^k \exists \vec{y}_i. \eta_i[\vec{x}_i, \vec{y}_i] \right), \quad (1)$$

where $\beta[\vec{x}]$ and $\eta_i[\vec{x}_i, \vec{y}_i]$ ($i = 1, \dots, k$) are conjunctions of atoms with variables \vec{x} and $\vec{x}_i \cup \vec{y}_i$, respectively. We call β *body* and $\bigvee_{i=1}^k \exists \vec{y}_i. \eta_i$ *head*. Bodies can be empty (we then omit \rightarrow), but heads must be non-empty. We require that \vec{x} and \vec{y}_i ($i = 1, \dots, k$) are mutually disjoint and that $\vec{x}_i \subseteq \vec{x}$ for all $i = 1, \dots, k$. We single out the fragment of *existential rules* by disallowing disjunction, i.e. requiring $k = 1$, and *Datalog rules* by disallowing existential quantifiers. We often omit the universal quantifiers from rules and treat conjunctions of atoms as sets of atoms.

Databases, Interpretations, and Entailment The semantics of formulas is based on logical interpretations, which we define as relational structures over a countably infinite set Nulls of *nulls*. A *schema* \mathcal{S} is a finite set of predicates. An *interpretation* \mathcal{I} over schema \mathcal{S} is a set of expressions $p(\vec{n})$ with $p \in \mathcal{S}$ and \vec{n} a list of nulls of length $Ar(p)$. We write Nulls(\mathcal{I}) for the set of nulls in \mathcal{I} . A *database* is a finite interpretation. See also the remarks on this notation below.

A *homomorphism* $h : \mathcal{I}_1 \rightarrow \mathcal{I}_2$ between interpretations \mathcal{I}_1 and \mathcal{I}_2 is a mapping h from the nulls in \mathcal{I}_1 to the nulls in \mathcal{I}_2 , such that $p(\vec{n}) \in \mathcal{I}_1$ implies $p(h(\vec{n})) \in \mathcal{I}_2$, where $h(\vec{n})$ denotes the list of h -values over \vec{n} . We also write $h(\mathcal{I})$ for $\{p(h(\vec{n})) \mid p(\vec{n}) \in \mathcal{I}\}$.

A *substitution* σ is a mapping from variables to nulls, which we extend to lists of variables and formulas as usual. A rule ρ as in (1) is *satisfied* by interpretation \mathcal{I} if every substitution $\sigma : \vec{x} \rightarrow \text{Nulls}$ with $\sigma(\beta) \subseteq \mathcal{I}$ can be extended to a substitution $\sigma' : \vec{x} \cup \vec{y}_i \rightarrow \text{Nulls}$ for some $i \in \{1, \dots, k\}$ such that $\sigma'(\eta_i) \subseteq \mathcal{I}$. Otherwise, if $\sigma(\beta) \subseteq \mathcal{I}$ but no extension σ' of σ verifies $\sigma'(\eta_i) \subseteq \mathcal{I}$ for some $i \in \{1, \dots, k\}$, then $\langle \rho, \sigma \rangle$ is *applicable* to \mathcal{I} .

\mathcal{I} *satisfies* a set Σ of rules if it satisfies every rule in Σ . An interpretation \mathcal{J} is *satisfied* by an interpretation \mathcal{I} if there is a homomorphism $h : \mathcal{J} \rightarrow \mathcal{I}$. \mathcal{I} is a *model* of a rule/rule set/interpretation/database \mathcal{X} if \mathcal{X} is satisfied by \mathcal{I} , written $\mathcal{I} \models \mathcal{X}$. As usual, we also write $\mathcal{X} \models \mathcal{Y}$ (“ \mathcal{X} entails \mathcal{Y} ”) if every model of \mathcal{X} is a model of \mathcal{Y} , where \mathcal{X} and \mathcal{Y} might be rules, rule sets, databases, or lists of several such elements. Note that the semantics of a database \mathcal{D} in this context corresponds to the semantics of a *Boolean conjunctive query* $\exists \vec{x}. \bigwedge \{p(x_{n_1}, \dots, x_{n_\ell}) \mid p(n_1, \dots, n_\ell) \in \mathcal{D}\}$ – we will therefore not introduce such queries as a separate notion. Also note that entailment and satisfaction between interpretations/databases coincide.

Abstract Queries, Expressivity, and Decidability An (*abstract*) *query* Ω over a schema \mathcal{S} is a set of databases over \mathcal{S} that is *closed under isomorphism*, i.e., such that whenever $\mathcal{D} \in \Omega$ and \mathcal{D}' is obtained from \mathcal{D} by bijective renaming of nulls, then $\mathcal{D}' \in \Omega$. Ω is further *closed under homomorphisms* if, for all $\mathcal{D} \in \Omega$ and homomorphisms $h : \mathcal{D} \rightarrow \mathcal{D}'$, we have $\mathcal{D}' \in \Omega$.

Definition 1. Let Goal be a nullary predicate. A query Ω over \mathcal{S} is *expressed* by a set Σ of rules if, for every database \mathcal{D} over \mathcal{S} , we have $\mathcal{D} \in \Omega$ if and only if $\Sigma, \mathcal{D} \models \text{Goal}$.

To discuss decidability of queries, we need to conceive databases as Turing machine inputs over a fixed alphabet. A *serialisation* for a schema \mathcal{S} is a word $s \in (\{0, 1, \|\} \cup \mathcal{S})^*$ of the form $e_1 \dots e_n$ where $n \geq 0$ and $e_i = p_i \| w_{i1} \| \dots \| w_{iAr(p_i)} \|$ for $w_{ij} \in \{0, 1\}^+$ and $p_i \in \mathcal{S}$. Given s of this form and an injection $\eta : \{0, 1\}^+ \rightarrow \text{Nulls}$, let $\eta(s)$ denote the database $\{p_i(\eta(w_{i1}), \dots, \eta(w_{iAr(p_i)})) \mid 1 \leq i \leq n\}$. Then s *corresponds* to a database \mathcal{D} if $\eta(s)$ is isomorphic to \mathcal{D} ; note that this does not depend on the choice of η .

A query Ω with schema \mathcal{S} is *decidable* if the set of all serialisations for \mathcal{S} that correspond to some $\mathcal{D} \in \Omega$ is a decidable language.

Remarks on our Notation Many works consider constants to appear in databases (not just nulls), but complexity and expressivity is usually studied for queries that are closed under isomorphisms, a.k.a. *generic* (Abiteboul, Hull, and Vianu 1995, Ch. 16), and nulls are more natural there. One can admit finitely many exceptions (elements that must not be renamed), but such “constants” can be simulated by marking them with dedicated unary predicates.

Specifying logical interpretations as sets of “atoms” that may use nulls is a notational convenience with some side effects: our interpretations cannot contain elements that do not stand in any relation, but they can have an empty domain. Both aspects do not change the notion of logical entailment *on the formulas we consider*.

Universal Models and the Chase Entailment of databases (corresponding to Boolean conjunctive queries) can be decided by considering only a subset of all models. Given sets \mathcal{J} and \mathcal{K} of interpretations, \mathcal{J} is *universal* for \mathcal{K} if, for all $\mathcal{K} \in \mathcal{K}$, there is $\mathcal{I} \in \mathcal{J}$ and a homomorphism $\mathcal{I} \rightarrow \mathcal{K}$. Consider a rule set Σ and database \mathcal{D} , and let \mathfrak{M} be the set of all models of Σ, \mathcal{D} . Then \mathcal{J} is a *universal model set* for Σ and \mathcal{D} if $\mathcal{J} \subseteq \mathfrak{M}$ and \mathcal{J} is universal for \mathfrak{M} .

Fact 1. *If \mathcal{J} is a universal model set for Σ and \mathcal{D} then, for every database \mathcal{C} , we have $\Sigma, \mathcal{D} \models \mathcal{C}$ iff $\mathcal{I} \models \mathcal{C}$ for all $\mathcal{I} \in \mathcal{J}$.*

Universal model sets can be computed with the *chase* algorithm. Here, we consider a variation of the *standard* (or *restricted*) chase for rules with disjunction, introduced by (Carral, Dragoste, and Krötzsch 2017).

Definition 2. A *chase tree* for a rule set Σ and database \mathcal{D} is a (finite or infinite) tree where each node is labelled by a database, such that:

1. The root is labelled with \mathcal{D} .
2. For every node with label \mathcal{E} that has ℓ children labelled $\mathcal{C}_1, \dots, \mathcal{C}_\ell$, there is a rule $\rho \in \Sigma$ of the form (1) and a substitution $\sigma : \vec{x} \rightarrow \text{Nulls}$ such that (i) $\langle \rho, \sigma \rangle$ is applicable to \mathcal{E} , (ii) ρ has $k = \ell$ head disjuncts, and (iii) $\mathcal{C}_i = \mathcal{E} \cup \sigma_i(\eta_i)$ where σ_i extends σ by mapping each variable $y \in \vec{y}_i$ to a fresh null.

3. For each rule $\rho \in \Sigma$ and each substitution σ , there is $i \geq 1$ such that $\langle \rho, \sigma \rangle$ is not applicable to the label of any node of depth $\geq i$.

The *result* that corresponds to a chase tree is the set of all interpretations that can be obtained as the union of all interpretations along a path in the tree.

Condition (3) ensures fair, exhaustive rule application, but different orders of application can lead to different chase trees, which can also have different results. Nevertheless, every result is semantically correct in the following sense:

Fact 2. *Every result of a chase on a rule set Σ and database \mathcal{D} is a universal model set for Σ and \mathcal{D} .*

The pair $\langle \Sigma, \mathcal{D} \rangle$ is *chase-terminating* if all its chase trees are finite – by König’s Lemma, this is equivalent to all chase results for $\langle \Sigma, \mathcal{D} \rangle$ containing only finite interpretations; this corresponds to *all-strategy termination*. Σ is *chase-terminating* if $\langle \Sigma, \mathcal{D} \rangle$ is chase-terminating for every database \mathcal{D} ; this corresponds to *universal termination*.

Turing Machines We will use (deterministic) Turing machines (TM), denoted as a tuple $M = \langle Q, \Gamma, \delta \rangle$, with states Q , tape alphabet Γ with blank $_ \in \Gamma$, and transition function δ . M has a distinguished initial state $q_S \in Q$, and accepting and rejecting halting states $q_A, q_R \in Q$. For all states $q \in Q \setminus \{q_A, q_R\}$ and tape symbols $a \in \Gamma$, there is exactly one transition $(q, a) \mapsto (r, b, D) \in \delta$. We assume that TM tapes are unbounded to the right but bounded to the left, and that TMs will never attempt to move left on the first position of the tape (this is w.l.o.g., since one can modify any TM to insert a marker at the tape start to recognise this case).

3 On the Expressivity of Disjunctive Rules

In this section, we show how to express a homomorphism-closed, decidable query with a disjunctive rule set. This construction will be the basis for finding a chase-terminating set of (deterministic) existential rules. Throughout this section, Ω is a fixed but arbitrary homomorphism-closed query over signature \mathcal{S} , and $M = \langle Q, \Gamma, \delta \rangle$ is a TM that decides Ω .

To express Ω , we specify five rule sets $\mathcal{R}_1 \subseteq \mathcal{R}_2 \subseteq \mathcal{R}_3 \subseteq \mathcal{R}_4 \subseteq \mathcal{R}_5$ (see Figures 1–6), which will be explained later. We want to show the following result:

Theorem 3. *The set \mathcal{R}_5 of disjunctive existential rules expresses the query Ω .*

To show this, we fix an arbitrary database \mathcal{D} over \mathcal{S} . Theorem 3 then follows from Fact 1 and the next lemma:

Lemma 4. *There is a universal model set \mathfrak{M} of \mathcal{R}_5 and \mathcal{D} such that $\mathcal{D} \in \Omega$ iff $\text{Goal} \in \mathcal{I}$ for every $\mathcal{I} \in \mathfrak{M}$.*

The universal model set \mathfrak{M} is a complicated structure that we describe step by step, by specifying five sets of interpretations – $\mathcal{J}_1, \mathcal{J}_2, \mathcal{J}_3, \mathcal{J}_4$, and \mathcal{J}_5 –, such that (1) \mathcal{J}_i is a universal model set of $\mathcal{R}_i, \mathcal{D}$ for each $1 \leq i \leq 5$, (2) $|\mathcal{J}_1| = |\mathcal{J}_2| = |\mathcal{J}_3| = |\mathcal{J}_4| = |\mathcal{J}_5|$, and (3) for each $1 \leq i < j \leq 5$ and each $\mathcal{I} \in \mathcal{J}_i$, there is exactly one $\mathcal{J} \in \mathcal{J}_j$ with $\mathcal{I} \subseteq \mathcal{J}$. Lemma 4 can then be shown using $\mathfrak{M} = \mathcal{J}_5$.

Before dwelling into the details of each rule set and universal model, we give an overview of the construction: \mathcal{R}_1

$$\rightarrow \exists y. \text{First}(y) \wedge \text{DbDom}(y) \quad (2)$$

$$\rightarrow \exists z. \text{Last}(z) \wedge \text{DbDom}(z) \quad (3)$$

$$\text{p}(\vec{x}) \rightarrow \text{In}_{\text{p}}(\vec{x}) \wedge \bigwedge_{x \in \vec{x}} \text{DbDom}(x) \quad (4)$$

$$\text{DbDom}(x) \rightarrow \text{Eq}(x, x) \quad (5)$$

$$\text{Eq}(x, y) \rightarrow \text{Eq}(y, x) \quad (6)$$

$$\text{NEq}(x, y) \rightarrow \text{NEq}(y, x) \quad (7)$$

$$\text{R}(\vec{x}) \wedge \text{Eq}(x_i, y) \rightarrow \text{R}(\vec{x}_{x_i \mapsto y}) \quad (8)$$

$$\text{DbDom}(x) \wedge \text{DbDom}(y) \rightarrow \text{Eq}(x, y) \vee \text{NEq}(x, y) \quad (9)$$

$$\text{LT}(x, y) \wedge \text{LT}(y, z) \rightarrow \text{LT}(x, z) \quad (10)$$

$$\text{First}(x) \wedge \text{NEq}(x, y) \rightarrow \text{LT}(x, y) \quad (11)$$

$$\text{NEq}(x, y) \wedge \text{Last}(y) \rightarrow \text{LT}(x, y) \quad (12)$$

$$\text{NEq}(x, y) \rightarrow \text{LT}(x, y) \vee \text{LT}(y, x) \quad (13)$$

$$\bigwedge_{x \in \vec{x}} \text{DbDom}(x) \rightarrow \text{In}_{\text{p}}(\vec{x}) \vee \text{NIn}_{\text{p}}(\vec{x}) \quad (14)$$

Figure 1: The rule set \mathcal{R}_1 , where rules (4) and (14) are instantiated for each $\text{p} \in \mathcal{S}$, and rules (8) are instantiated for each $\text{R} \in \{\text{First}, \text{Last}, \text{Eq}, \text{NEq}, \text{LT}\} \cup \{\text{In}_{\text{p}}, \text{NIn}_{\text{p}} \mid \text{p} \in \mathcal{S}\}$ and $1 \leq i \leq \text{Ar}(\text{R})$, and $\vec{x}_{x_i \mapsto y}$ denotes \vec{x} with x_i replaced by y .

$$\text{First}(x) \rightarrow \exists u. \text{Root}(u) \wedge \text{Rep}(x, u) \quad (15)$$

$$\text{Rep}(x, v) \wedge \text{LT}(x, z) \rightarrow \exists w. \text{Chi}(v, w) \wedge \text{Rep}(z, w) \quad (16)$$

$$\text{Last}(x) \wedge \text{Rep}(x, u) \rightarrow \text{Leaf}(u) \quad (17)$$

$$\text{Rep}(x, u) \wedge \text{Eq}(x, y) \rightarrow \text{Rep}(y, u) \quad (18)$$

$$\text{In}_{\text{p}}(\vec{x}) \wedge \bigwedge_{i=1}^{|\vec{x}|} \text{Rep}(x_i, u_i) \rightarrow \text{In}_{\text{p}}'(\vec{u}) \quad (19)$$

$$\text{NIn}_{\text{p}}(\vec{x}) \wedge \bigwedge_{i=1}^{|\vec{x}|} \text{Rep}(x_i, u_i) \rightarrow \text{NIn}_{\text{p}}'(\vec{u}) \quad (20)$$

Figure 2: The rule set \mathcal{R}_2 contains \mathcal{R}_1 (see Figure 1) and all above rules, where (19) and (20) are instantiated for each $\text{p} \in \mathcal{S}$.

constructs all possible linear orders over the nulls in \mathcal{D} , as well as all possible completions of \mathcal{D} with facts built using these nulls; $\mathcal{R}_2 \setminus \mathcal{R}_1$ extracts successor relations from the linear orders; $\mathcal{R}_3 \setminus \mathcal{R}_2$ associates to nulls representations of their positions in successor relations; $\mathcal{R}_4 \setminus \mathcal{R}_3$ encodes all initial TM configurations corresponding to some linear order and completion, and $\mathcal{R}_5 \setminus \mathcal{R}_4$ simulates the run of the TM on these configurations.

\mathcal{R}_1 : Linear Order and Database Completion \mathcal{R}_1 serves two distinct purposes: (1) predicates First , Last , Eq (“=”), NEq (“≠”), and LT (“<”) encode representations of possible linear orders over nulls in \mathcal{D} (collected in predicate DbDom); and (2) predicates In_{p} and NIn_{p} for each $\text{p} \in \mathcal{S}$ explicitly encode positive and negative (absent) facts in \mathcal{D} . Both purposes require disjunctive reasoning. Possible models include representations of strict, total linear orders (1) and the exact database completion (2), but also models for collapsed orders and inconsistent completions. The latter is not problematic since we consider homomorphism-closed queries.

We define the interpretation set \mathcal{J}_1 on the set $\Delta = \text{Nulls}(\mathcal{D}) \cup \{u_\alpha, u_\omega\}$, for fresh nulls $u_\alpha, u_\omega \notin \text{Nulls}(\mathcal{D})$. A (partially collapsed) linear order can be represented by

$$\text{Root}(u) \rightarrow \exists y_1, y_2. \text{Enc}(u, y_1, y_2) \wedge \text{S}_0(y_1) \wedge \text{Nxt}(y_1, y_2) \wedge \text{S}_1(y_2) \quad (21)$$

$$\text{Enc}(u, y_1, y_-) \wedge \text{Chi}(u, v) \rightarrow \exists z_1, z_-. \text{Enc}(v, z_1, z_-) \wedge \text{Cpy}_{+1}(y_1, y_-, z_1, z_-) \quad (22)$$

$$\text{Cpy}_{+1}(y_1, y_2, z_1, z_-) \wedge \text{S}_0(y_1) \wedge \text{Nxt}(y_1, y_2) \rightarrow \text{S}_1(z_1) \wedge \text{Nxt}(z_1, z_-) \wedge \text{S}_1(z_-) \quad (23)$$

$$\text{Cpy}_{+1}(y_1, y_2, z_1, z_-) \wedge \text{S}_1(y_1) \wedge \text{Nxt}(y_1, y_2) \rightarrow \exists z_2. \text{S}_0(z_1) \wedge \text{Nxt}(z_1, z_2) \wedge \text{S}_0(z_2) \wedge \text{Nxt}(z_2, z_-) \wedge \text{S}_1(z_-) \quad (24)$$

$$\text{Cpy}_{+1}(y_1, y_-, z_1, z_-) \wedge \text{S}_0(y_1) \wedge \text{Nxt}(y_1, y_2) \wedge \text{Nxt}(y_2, y_3) \rightarrow \exists z_2. \text{Cpy}(y_2, y_-, z_2, z_-) \wedge \text{S}_1(z_1) \wedge \text{Nxt}(z_1, z_2) \quad (25)$$

$$\text{Cpy}_{+1}(y_1, y_-, z_1, z_-) \wedge \text{S}_1(y_1) \wedge \text{Nxt}(y_1, y_2) \wedge \text{Nxt}(y_2, y_3) \rightarrow \exists z_2. \text{Cpy}_{+1}(y_2, y_-, z_2, z_-) \wedge \text{S}_0(z_1) \wedge \text{Nxt}(z_1, z_2) \quad (26)$$

$$\text{Cpy}(y_1, y_2, z_1, z_2) \wedge \text{S}_*(y_1) \wedge \text{Nxt}(y_1, y_2) \rightarrow \text{S}_*(z_1) \wedge \text{Nxt}(z_1, z_2) \wedge \text{S}_1(z_2) \quad (27)$$

$$\text{Cpy}(y_1, y_-, z_1, z_-) \wedge \text{S}_*(y_1) \wedge \text{Nxt}(y_1, y_2) \wedge \text{Nxt}(y_2, y_3) \rightarrow \exists z_2. \text{Cpy}(y_2, y_-, z_2, z_-) \wedge \text{S}_*(z_1) \wedge \text{Nxt}(z_1, z_2) \quad (28)$$

Figure 3: The rule set \mathcal{R}_3 contains \mathcal{R}_2 (Figure 2) and all of the above rules, where (27) and (28) are instantiated for each $* \in \{0, 1\}$.

an ordered partition $\vec{T} = T_1, \dots, T_k$ ($k \geq 1$) of Δ where $u_\alpha \in T_1$ and $u_\omega \in T_k$. Let $Ords$ be the set of all such \vec{T} , and let $Compls(\vec{T})$ be the set of all interpretations with nulls Δ that are set-minimal among the models of \mathcal{R}_1 and \mathcal{D} , and that contain the database

$$\{\text{First}(u) \mid u \in T_1\} \cup \{\text{Last}(u) \mid u \in T_k\} \cup \\ \{\text{Eq}(t, u) \mid 1 \leq i \leq k; t, u \in T_i\} \cup \\ \{\text{LT}(t, u), \text{NEq}(t, u), \text{NEq}(u, t) \mid 1 \leq i < j \leq k; t \in T_i; u \in T_j\}.$$

By minimality, each $\mathcal{I} \in Compls(\vec{T})$ contains exactly one of $\{\text{In}_p(\vec{t}), \text{NIn}_p(\vec{t})\}$ for every $\vec{t} \subseteq \Delta$ with $|\vec{t}| = Ar(p)$.

Lemma 5. $\mathfrak{J}_1 = \bigcup_{\vec{T} \in Ords} Compls(\vec{T})$ is a universal model set of \mathcal{R}_1 and \mathcal{D} .

Proof Sketch. We can construct a chase tree for \mathcal{R}_1 and \mathcal{D} that prioritises the application of rules in the order of their appearance in Figure 1. The result of that chase \mathfrak{K} is a finite universal model set (Fact 2) and we can show that every $\mathcal{K} \in \mathfrak{K}$ is isomorphic to a unique $\mathcal{I} \in \mathfrak{J}_1$. \square

Each set \mathfrak{J}_i for some $2 \leq i \leq 5$ is obtained as $\mathfrak{J}_i = \{\text{Int}_i(\mathcal{I}) \mid \mathcal{I} \in \mathfrak{J}_{i-1}\}$ for a function Int_i as defined below. Every $\mathcal{I} \in \mathfrak{J}_i$ with $1 \leq i \leq 5$ contains a unique interpretation $\text{Seed}(\mathcal{I}) \in \mathfrak{J}_1$, and there is a unique ordered partition $\text{Order}(\mathcal{I}) \in Ords$ such that $\text{Seed}(\mathcal{I}) \in Compls(\text{Order}(\mathcal{I}))$.

\mathcal{R}_2 : Representative Tree The purpose of \mathcal{R}_2 is to extract successor relations from the transitive linear order LT. Given an ordered partition $\text{Order}(\mathcal{I}) = T_1, \dots, T_k$ of some model \mathcal{I} of \mathcal{R}_1 , \mathcal{R}_2 constructs a finite tree structure – defined using predicates Root , Chi (“child”), and Leaf – where each path represents a sub-sequence T_{z_1}, \dots, T_{z_p} of T_1, \dots, T_k with $z_1 = 1$ and $z_p = k$. Unavoidably, some paths will skip some T_i , but it suffices for our purposes if one path is complete. The elements of any set T_j are related to the tree nodes that represent T_j by a predicate Rep . Moreover, nodes are related via predicates In'_p and NIn'_p that reflect the relations for In_p and NIn_p that hold between the represented elements (in the database completion of the considered model \mathcal{I} of \mathcal{R}_1).

Given \mathcal{I} with $|\text{Order}(\mathcal{I})| = k$ as above, let \mathbf{Z} be the set of all words $z_1 \dots z_p \in \{1, \dots, k\}^*$ such that $z_1 = 1$ and $z_i < z_{i+1}$ for all $i \in \{1, \dots, p-1\}$. Moreover, let

$\text{end}(z_1 \dots z_p) := z_p$. Using fresh nulls $\{u_w \mid w \in \mathbf{Z}\}$, we define $\text{Tree}(\mathcal{I})$ to be the interpretation

$$\{\text{Root}(u_1)\} \cup \{\text{Chi}(u_w, u_{wz}) \mid wz \in \mathbf{Z}, 2 \leq z \leq k\} \cup \\ \{\text{Leaf}(u_w) \mid w \in \mathbf{Z}; \text{end}(w) = k\} \cup \\ \{\text{Rep}(x, u_w) \mid w \in \mathbf{Z}; x \in T_{\text{end}(w)}\}.$$

Now $\text{Int}_2(\mathcal{I})$ is the least interpretation that contains \mathcal{I} and $\text{Tree}(\mathcal{I})$, and that further satisfies rules (19) and (20). We can extend Lemma 5 as follows. The required universal model set is obtained by any chase that prioritises rule (18).

Lemma 6. $\mathfrak{J}_2 = \{\text{Int}_2(\mathcal{I}) \mid \mathcal{I} \in \mathfrak{J}_1\}$ is a universal model set of \mathcal{R}_2 and \mathcal{D} .

\mathcal{R}_3 : Position Binary Encodings The purpose of \mathcal{R}_3 is to associate each node in the tree of \mathcal{R}_2 with a binary encoding of its distance from the root (the root starts with “distance” 2 for technical reasons). Encodings start at the least significant bit and always end in 1 (i.e., have no leading 0s). To simplify upcoming steps, encodings take the form of little TM tapes, represented by a Nxt -connected chain of nulls with unary predicates S_0 and S_1 encoding the symbol at each position. Nodes u relate to the first and last null t_s and t_e of their “tape” through facts $\text{Enc}(u, t_s, t_e)$. Facts $\text{Cpy}(a_s, a_e, b_s, b_e)$ are used to create a tape between b_s and b_e that contains a copy of the information on the tape between a_s and a_e . Predicate Cpy_{+1} is analogous, but creates a representation of the successor of the number that is copied.

Consider a model \mathcal{I} of \mathcal{R}_2 and define the set of sequences \mathbf{Z} as before. For $w \in \mathbf{Z}$ of length $|w|$, and $b_1 \dots b_\ell$ the binary representation of $|w| + 1$, let $\text{EncPos}(w)$ be the database

$$\{\text{Enc}(u_w, e_w^1, e_w^\ell)\} \cup \{\text{S}_{b_i}(e_w^i) \mid 1 \leq i \leq \ell\} \cup \\ \{\text{Nxt}(e_w^{i-1}, e_w^i) \mid 2 \leq i \leq \ell\}.$$

Let $\mathcal{J} = \mathcal{I} \cup \bigcup_{w \in \mathbf{Z}} \text{EncPos}(w)$. We define $\text{Int}_3(\mathcal{I})$ as the smallest superset of \mathcal{J} that satisfies all rules in \mathcal{R}_3 while including only the nulls in \mathcal{J} . $\text{Int}_3(\mathcal{I})$ extends \mathcal{J} only by missing Cpy and Cpy_{+1} relations, which can be inferred by slightly rewritten rules. For example, rule (22) is satisfied when applying the following rule to \mathcal{J} :

$$\text{Enc}(u, y_1, y_-) \wedge \text{Chi}(u, v) \wedge \text{Enc}(v, z_1, z_-) \\ \rightarrow \text{Cpy}_{+1}(y_1, y_-, z_1, z_-).$$

$$\text{Ld}_p(u, t, \vec{v}) \wedge \text{In}'_p(\vec{v}) \rightarrow \exists \vec{x}, y, \text{Sp}(t) \wedge \text{Nxt}(t, x_1) \wedge \bigwedge_{i=1}^{|\vec{v}|} \text{LdE}(v_i, x_i, x_{i+1}) \wedge \text{Nxt}(x_{|\vec{v}|+1}, y) \wedge \text{Rdy}_p(u, y, \vec{v}) \quad (29)$$

$$\text{LdE}(v, x_s, x_e) \wedge \text{Enc}(v, y_1, y_e) \rightarrow \exists z_1, z_e, \text{S}_{||}(x_s) \wedge \text{Nxt}(x_s, z_1) \wedge \text{Cpy}(y_1, y_e, z_1, z_e) \wedge \text{Nxt}(z_e, x_e) \wedge \text{S}_{||}(x_e) \quad (30)$$

Figure 4: The rule set \mathcal{R}_4 contains \mathcal{R}_3 (see Figure 3), the rules from Figure 5, and the above rules instantiated for all $p \in \mathcal{S}$.

$$\text{Leaf}(u) \rightarrow \exists t, \text{Ld}_1(u, t, u) \wedge \text{Hd}_{q_s}(t) \quad (31)$$

$$\text{Ld}_\ell(u, t, \vec{v}) \rightarrow \text{Ld}_{p_1^\ell}(u, t, \vec{v}) \quad (32)$$

$$\text{Rdy}_{p_j^\ell}(u, t, \vec{v}) \rightarrow \text{Ld}_{p_{j+1}^\ell}(u, t, \vec{v}) \quad (33)$$

$$\text{Rdy}_{p_{\bar{n}}^\ell}(u, t, \vec{v}) \rightarrow \text{Rdy}_\ell(u, t, \vec{v}) \quad (34)$$

$$\begin{aligned} \text{Rdy}_\ell(u, t, \vec{v}) \wedge \bigwedge_{i=k+1}^\ell \text{Root}(v_i) \wedge \text{Chi}(w, v_k) \\ \rightarrow \text{Ld}_\ell(u, t, v_1, \dots, v_{k-1}, w, u, \dots, u) \end{aligned} \quad (35)$$

$$\text{Rdy}_\ell(u, t, \vec{v}) \wedge \bigwedge_{i=1}^\ell \text{Root}(v_i) \rightarrow \text{Ld}_{\ell+1}(u, t, u, \dots, u) \quad (36)$$

$$\text{Rdy}_{\bar{m}}(u, t, \vec{v}) \wedge \bigwedge_{i=1}^{\bar{m}} \text{Root}(v_i) \rightarrow \text{S}_-(t) \wedge \text{End}(t) \quad (37)$$

$$\text{Ld}_p(u, t, \vec{v}) \wedge \text{NI}'_p(\vec{v}) \rightarrow \text{Rdy}_p(u, t, \vec{v}) \quad (38)$$

Figure 5: Some rules of \mathcal{R}_4 , to be instantiated for all $1 \leq j \leq \bar{n}-1$, $1 \leq k \leq \ell \leq \bar{m}$, and $p \in \mathcal{S}$.

All other rules can be rewritten analogously, since every existentially quantified variable is used in unique ways with predicates other than Cpy and Cpy_{+1} .

For every $\mathcal{I} \in \mathcal{I}_2$, we show that $\text{Int}_3(\mathcal{I})$ is isomorphic to a result of the chase on \mathcal{R}_3 and \mathcal{I} . The next result then follows from Lemma 6.

Lemma 7. $\mathcal{I}_3 = \{\text{Int}_3(\mathcal{I}) \mid \mathcal{I} \in \mathcal{I}_2\}$ is a universal model set of \mathcal{R}_3 and \mathcal{D} .

\mathcal{R}_4 : Initial TM Configuration For each leaf in the tree of completions, \mathcal{R}_4 creates the representation of an initial TM configuration. The tape is again represented by a Nxt -chain, using further unary predicates $\text{S}_{||}$, S_- , and S_p (for all $p \in \mathcal{S}$) for additional tape symbols. Hd_{q_s} marks the TM's starting position and initial state q_s , and End the end of the tape.

Let \bar{m} be the maximal arity of predicates in \mathcal{S} . We require that there is some $\bar{n} > 0$ such that \mathcal{S} contains exactly \bar{n} predicates $\text{p}_1^i, \dots, \text{p}_{\bar{n}}^i$ of arity i , for every $1 \leq i \leq \bar{m}$. This is without loss of generality, except for the exclusion of nullary predicates. Our results do not depend on this restriction, but it helps to simplify the presentation of our main ideas.

To serialise the data as a tape, we iterate over all predicate arities $\ell = 1, \dots, \bar{m}$ and over all lists \vec{v} of tree nodes with length ℓ . In this process, $\text{Ld}_\ell(u, t, \vec{v})$ expresses that, while encoding the leaf u , after constructing the tape until position t , we continue serialising ℓ -ary predicate data for arguments \vec{v} . Analogously, $\text{Rdy}_\ell(u, t, \vec{v})$ means that this was completed at tape position t . Similar predicates Ld_p and Rdy_p are used to consider a specific predicate $p \in \mathcal{S}$ during this process. The rules in Figure 5 start the serialisation (31), proceed over all predicates (32)–(34), iterate over parameter vectors (35) and arities (36), and finally end the tape (37).

Absent facts do not need to be serialised (38), while present facts can be treated by copying the encodings for

each of their parameters (29) and (30). In the latter, LdE states that a specific argument is serialised between two given tape positions.

The resulting TM tapes serialise facts $\text{In}'_p(\vec{u})$ as introduced by \mathcal{R}_2 , i.e., where \vec{u} are nodes in the representative tree. Given a model $\mathcal{I} \in \mathcal{I}_3$ with some $\text{Leaf}(u_w) \in \mathcal{I}$, let $\text{branch}(u_w)$ be the set of all nodes $u_{w'}$ on the branch of u_w , i.e. all nulls $u_{w'}$ where w' is a prefix of w . Elements of $\text{branch}(u_w)$ are totally ordered by setting $u_{w_1} \prec u_{w_2}$ if $|w_1| > |w_2|$. Predicates are totally ordered by setting $\text{p}_a^i \prec \text{p}_b^j$ if either $a < b$, or both $a = b$ and $i < j$. We can then order facts as $\text{p}(\vec{u}) \prec \text{q}(\vec{v})$ if $\vec{u}, \vec{v} \subseteq \text{branch}(u_w)$ and $\langle \text{p}, \vec{u} \rangle$ is lexicographically before $\langle \text{q}, \vec{v} \rangle$.

Now let $\text{branchDb}(\mathcal{I}, u_w) = \{\text{p}(\vec{u}) \mid \text{In}'_p(\vec{u}) \in \mathcal{I}, \vec{u} \subseteq \text{branch}(u_w)\}$ denote the set of all facts on the branch with leaf u_w , and let $\text{branchTape}(\mathcal{I}, u_w)$ denote the TM tape serialisation (as defined in Section 2) of $\text{branchDb}(\mathcal{I}, u_w)$ according to the total order \prec and representing each node u_w by the binary representation of $|w| + 1$ as before. Given $S = \text{branchTape}(\mathcal{I}, u_w)$, let $\text{startConf}(\mathcal{I}, u_w)$ be the following interpretation:

$$\{\text{Ld}_1(u_w, t_w^1, u_w), \text{Hd}_{q_s}(t_w^1), \text{End}(t_w^{|S|+1})\} \cup$$

$$\{\text{Nxt}(t_w^{j-1}, t_w^j) \mid 2 \leq j \leq |S| + 1\} \cup$$

$$\{\text{S}_a(t_w^j) \mid 1 \leq j \leq |S|, a = S[j]\} \cup \{\text{S}_-(t_w^{|S|+1})\}.$$

Let \mathcal{J} be the extension of \mathcal{I} with $\text{startConf}(\mathcal{I}, u_w)$ for every $\text{Leaf}(u_w) \in \mathcal{I}$. We define $\text{Int}_4(\mathcal{I})$ to be the smallest superset of \mathcal{J} that satisfies all rules in \mathcal{R}_4 while including only the nulls in \mathcal{J} . As in the case of Int_3 , the missing relations can easily be inferred using the original rules or, for (29) and (30), with simple rewritings thereof.

Lemma 8. $\mathcal{I}_4 = \{\text{Int}_4(\mathcal{I}) \mid \mathcal{I} \in \mathcal{I}_3\}$ is a universal model set of \mathcal{R}_4 and \mathcal{D} .

\mathcal{R}_5 : TM Run The purpose of \mathcal{R}_5 is to simulate the run of the deterministic TM $\langle Q, \Gamma, \delta \rangle$ on each of the initial tapes created by \mathcal{R}_4 . We continue to use predicate Nxt for neighbouring tape cells (augmented with its transitive closure Nxt^+), S_b to encode tape symbols $b \in \Gamma$, and Hd_q to encode head position and current state $q \in Q$. Predicate Stp connects tape cells in each configuration to the corresponding tape cells in the next configuration (provided the TM performs another step). The rules in Figure 6 are a standard TM encoding, with the slight exception of rule (43), which adds a new blank tape cell in each step (even if not used by the TM). Our rules use the assumptions on TMs in Section 2.

Consider some $\mathcal{I} \in \mathcal{I}_4$. It is easy and only mildly laborious to define interpretations $\text{Run}(u_w)$ that represent all successor configurations of the starting configuration $\text{startConf}(\mathcal{I}, u_w)$, appropriately connected with Stp and

$$\text{Hd}_{q_A}(x) \rightarrow \text{Goal} \quad (39)$$

$$\text{Nxt}(x, y) \rightarrow \text{Nxt}^+(x, y) \quad (40)$$

$$\text{Nxt}^+(x, y) \wedge \text{Nxt}^+(y, z) \rightarrow \text{Nxt}^+(x, z) \quad (41)$$

$$\text{Nxt}(x, y) \wedge \text{Stp}(x, z) \wedge \text{Stp}(y, w) \rightarrow \text{Nxt}(z, w) \quad (42)$$

$$\text{End}(x) \wedge \text{Stp}(x, z) \rightarrow \exists v. \text{Nxt}(z, v) \wedge \text{S}_-(v) \wedge \text{End}(v) \quad (43)$$

$$\text{Hd}_q(x) \wedge \text{S}_a(x) \rightarrow \exists z. \text{Stp}(x, z) \wedge \text{S}_b(z) \quad (44)$$

$$\text{Hd}_q(x) \wedge \text{Nxt}^+(x, y) \wedge \text{S}_c(y) \rightarrow \exists z. \text{Stp}(y, z) \wedge \text{S}_c(z) \quad (45)$$

$$\text{Hd}_q(x) \wedge \text{Nxt}^+(y, x) \wedge \text{S}_c(y) \rightarrow \exists z. \text{Stp}(y, z) \wedge \text{S}_c(z) \quad (46)$$

$$\text{Hd}_q(x) \wedge \text{S}_a(x) \wedge \text{Stp}(x, z) \wedge \text{Nxt}(z, w) \rightarrow \text{Hd}_r(w) \quad (47)$$

$$\text{Hd}_q(x) \wedge \text{S}_a(x) \wedge \text{Stp}(x, z) \wedge \text{Nxt}(w, z) \rightarrow \text{Hd}_r(w) \quad (48)$$

Figure 6: The rule set \mathcal{R}_5 contains \mathcal{R}_4 (see Figure 4) and the above rules, where we instantiate rules (44)–(46) for all transitions $(q, a) \mapsto (r, b, X) \in \delta$ and $c \in \Gamma$; rule (47) for all $(q, a) \mapsto (r, b, +1) \in \delta$; and rule (48) for all $(q, a) \mapsto (r, b, -1) \in \delta$.

the transitive closure Nxt^+ . Moreover, let \mathcal{J} be the extension of \mathcal{I} with all Nxt^+ required to satisfy (40) and (41) (note that Nxt also occurs in encodings from \mathcal{R}_3). We define $\text{Int}_5(\mathcal{I})$ as the union of \mathcal{J} with the interpretations $\text{Run}(u_w)$ for all u_w with $\text{Leaf}(u_w) \in \mathcal{I}$.

Lemma 9. $\mathfrak{J}_5 = \{\text{Int}_5(\mathcal{I}) \mid \mathcal{I} \in \mathfrak{J}_4\}$ is a universal model set of \mathcal{R}_5 and \mathcal{D} .

Proving Lemma 4 To complete the proof of Lemma 4, we set $\mathfrak{M} = \mathfrak{J}_5$. For $\mathcal{I} \in \mathfrak{M}$, let $\text{Db}(\mathcal{I}) = \{\mathfrak{p}(\vec{t}) \mid \text{In}_{\mathfrak{p}}(\vec{t}) \in \mathcal{I}\}$ denote the completed database created by \mathcal{R}_1 . Due to rule (4) in Figure 1, there is a homomorphism $\mathcal{D} \rightarrow \text{Db}(\mathcal{I})$. Moreover, the representation tree constructed for \mathcal{I} by \mathcal{R}_2 has a branch that is maximal, i.e., has $|\text{Order}(\mathcal{I})|$ nodes; this branch has a leaf u_w with $|w| = |\text{Order}(\mathcal{I})|$. We obtain a homomorphism $\text{Db}(\mathcal{I}) \rightarrow \text{branchDb}(\mathcal{I}, u_w)$. Lemma 4 now follows from Lemma 9 and Lemmas 11 and 10 below.

Lemma 10. If $\mathcal{D} \in \mathfrak{Q}$, then $\text{Goal} \in \mathcal{I}$ for each $\mathcal{I} \in \mathfrak{M}$.

Proof. As shown above, there is a homomorphism $\mathcal{D} \rightarrow \text{branchDb}(\mathcal{I}, u_w)$ for the node u_w where $|w| = |\text{Order}(\mathcal{I})|$. Since \mathfrak{Q} is closed under homomorphisms, $\mathcal{D} \in \mathfrak{Q}$ implies $\text{branchDb}(\mathcal{I}, u_w) \in \mathfrak{Q}$. By the correctness of our TM simulation, we obtain $\text{Goal} \in \mathcal{I}$. \square

Lemma 11. If $\mathcal{D} \notin \mathfrak{Q}$, then $\text{Goal} \notin \mathcal{I}$ for some $\mathcal{I} \in \mathfrak{M}$.

Proof. Consider some $\mathcal{I} \in \mathfrak{M}$ such that $\text{Db}(\mathcal{I}) = \mathcal{D}$ and $\text{NEq}(t, u) \in \mathcal{I}$ for each $t, u \in \text{Nulls}(\mathcal{D})$ with $t \neq u$. Let u_w denote the leaf node with $|w| = |\text{Order}(\mathcal{I})|$ as before. Then $\text{branchDb}(\mathcal{I}, u_w)$ is isomorphic to $\text{Db}(\mathcal{I}) = \mathcal{D}$. By the correctness of our TM simulation, Goal is not derived from this maximal branch. Moreover, for all other leaf nodes u_v with $\text{Leaf}(u_v) \in \mathcal{I}$, there is a homomorphism $\text{branchDb}(\mathcal{I}, u_v) \rightarrow \text{branchDb}(\mathcal{I}, u_w)$. Since \mathfrak{Q} is closed under homomorphisms, the TM does not accept any such $\text{branchDb}(\mathcal{I}, u_v)$, so $\text{Goal} \notin \mathcal{I}$. \square

4 Ensuring Chase Termination

While the rules in Section 3 are semantically correct, the disjunctive chase may not terminate on them. Many known fragments of existential rules can guarantee chase termination, including for expressive cases where termination might be exponential (Carral et al. 2019), but they are not applicable to our case, since the runtime of TMs that decide a query can in general not be bounded by any elementary function. Indeed, we rely on the TM to stop “naturally”, by virtue of being a decider. Nevertheless, our rules lead to infinite chase trees, e.g., if the disjunctive guessing of LT leads to a cycle, which enables rule (16) to create an infinite path in the representation tree. We will now show that this can be avoided:

Theorem 12. *Every homomorphism-closed decidable query is expressed by a set of disjunctive rules that is chase-terminating for all databases over the schema of the query.*

To show this, we refine and generalise the “emergency brake” technique of Krötzsch, Marx, and Rudolph (2019), and re-formulate it as a general rule set transformation. This not only yields a generic method that is of independent interest, but it also allows us to address potential termination problems in our prior modelling.

Definition 3. Consider a rule set Σ and a nullary predicate Halt that does not occur in Σ . For every predicate \mathfrak{p} in Σ , let $\hat{\mathfrak{p}}$ be a fresh predicate of the same arity, and, for any formula ψ , let $\hat{\psi}$ be ψ with all predicates \mathfrak{p} replaced by $\hat{\mathfrak{p}}$. Now the set $\text{brake}(\Sigma, \text{Halt})$ consists of the following rules:

$$\rightarrow \exists v. \text{Brake}(v) \quad (49)$$

$$\text{Halt} \wedge \text{Brake}(x) \rightarrow \text{Real}(x) \quad (50)$$

$$\hat{\mathfrak{p}}(\vec{x}) \wedge \bigwedge_{x \in \vec{x}} \text{Real}(x) \rightarrow \mathfrak{p}(\vec{x}) \quad \text{for all } \mathfrak{p} \text{ in } \Sigma \quad (51)$$

For every rule $\rho : \beta[\vec{x}] \rightarrow \bigvee_{i=1}^k \exists \vec{y}_i. \eta_i[\vec{x}_i, \vec{y}_i]$:

$$\beta[\vec{x}] \wedge \text{Brake}(v) \rightarrow \bigvee_{i=1}^k (\mathfrak{B}_\rho^i(\vec{x}_i) \wedge \hat{\eta}_i[\vec{x}_i, \vec{y}_i \mapsto v]) \wedge \bigwedge_{x \in \vec{x}_i} \text{Real}(x) \quad (52)$$

$$\mathfrak{B}_\rho^i(\vec{x}_i) \rightarrow \exists \vec{y}_i. \hat{\eta}_i[\vec{x}_i, \vec{y}_i] \wedge \bigwedge_{y \in \vec{y}_i} \text{Real}(y) \quad (53)$$

where $\hat{\eta}_i[\vec{x}_i, \vec{y}_i \mapsto v]$ is $\hat{\eta}_i$ with each variable $y \in \vec{y}_i$ replaced by v , and Brake , Real , and all \mathfrak{B}_ρ^i are fresh predicates with arities as indicated.

Note that $\text{brake}(\Sigma, \text{Halt})$ does not define rules to derive Halt , and indeed the transformation largely preserves the models of Σ in the following sense:

Lemma 13. *Consider a rule set Σ and database \mathcal{D} over predicates that occur in Σ . For every model \mathcal{I} of $\text{brake}(\Sigma, \text{Halt})$ and \mathcal{D} , the set $\mathcal{I}^- = \{\mathfrak{p}(\vec{n}) \mid \mathfrak{p}(\vec{n}) \in \mathcal{I}, \mathfrak{p} \text{ occurs in } \Sigma\}$ is a model of Σ and \mathcal{D} , and every model \mathcal{J} of Σ and \mathcal{D} is of this form.*

Proof. Consider a rule $\rho \in \Sigma$ as in Definition 3, and let σ be a substitution such that $\sigma(\beta) \subseteq \mathcal{I}^-$. Then we can apply rules (52), (53), and finally (51) to derive $\sigma'(\eta_i) \subseteq \mathcal{I}^-$ for a suitable extension σ' of σ . Hence $\mathcal{I}^- \models \Sigma$.

Conversely, let $\mathcal{J} \models \Sigma$. A model \mathcal{I} of $\text{brake}(\Sigma, \text{Halt})$ can be found by adding, for each matching body $\sigma(\beta) \subseteq \mathcal{J}$

of rule ρ , an atom $\sigma(B_\rho^i(\vec{x}_i))$ for some i such that $\sigma'(\eta_i) \subseteq \mathcal{J}$ for an extension σ' of σ . To obtain the required model \mathcal{I} of $\text{brake}(\Sigma, \text{Halt})$, it remains to add facts $\text{Brake}(b)$ for a fresh null b , $\sigma(\hat{\eta}_i[\vec{x}_i, \vec{y}_i \mapsto b])$ as in (52) for every $\sigma(B_\rho^i(\vec{x}_i)) \in \mathcal{I}$, and $\text{Real}(n)$ for every $p(\vec{n}) \in \mathcal{J}$ and $n \in \vec{n}$. \square

For $\text{brake}(\Sigma, \text{Halt})$ to be useful, we need to add rules that can “pull the brake” by deriving Halt . Doing so stops the chase in the following sense:

Lemma 14. *Consider a rule set Σ , a database \mathcal{D} over predicates that occur in Σ , and a set Π of rules of the form $\beta \rightarrow \text{Halt}$ where β only uses predicates in Σ . If \mathcal{I} is the label of a node in a chase tree for $\Sigma \cup \Pi$ and \mathcal{D} such that $\text{Halt} \in \mathcal{I}$, then the tree starting at the node of \mathcal{I} is finite.*

Proof. Since $\text{Halt} \in \mathcal{I}$, there is a substitution σ such that $\langle \rho_{(50)}, \sigma \rangle$ is applicable (for $\rho_{(50)}$ in (50)). By fairness, $\text{Real}(\sigma(x))$ will be derived at some depth of the tree. From this depth on, no rule of form (53) is applicable: given $\text{Real}(\sigma(x))$, the head of rules of form (52) already satisfies the head of the rule (53) that could be applied to a newly derived atom for B_ρ^i . Rules other than (53) do not contain existential quantifiers thus can only be applied a finite number of times before the chase on this part of the tree terminates. \square

If Halt is derived, the semantic correspondence of Lemma 13 is weakened, but suffices to preserve entailments:

Lemma 15. *Consider Σ , \mathcal{D} , and Π as in Lemma 14. For every model \mathcal{I} of $\text{brake}(\Sigma, \text{Halt}) \cup \Pi$ and \mathcal{D} , \mathcal{I}^- (as in Lemma 13) is a model of Σ and \mathcal{D} .*

Proof. This is immediate from Lemma 13 and the fact that every model of $\text{brake}(\Sigma, \text{Halt}) \cup \Pi$ and \mathcal{D} is also a model of $\text{brake}(\Sigma, \text{Halt})$ and \mathcal{D} . \square

Having established the key properties of the emergency brake construction, we can now apply it to show Theorem 12. Given the rule set \mathcal{R}_5 as defined for a query Ω in Section 3, let \mathcal{R}_6 denote the extension of $\text{brake}(\mathcal{R}_5, \text{Halt})$ with the following rules:

$$\text{In}_p(\vec{x}) \wedge \text{NIn}_p(\vec{x}) \rightarrow \text{Halt} \quad (54)$$

$$\text{LT}(x, x) \rightarrow \text{Halt} \quad (55)$$

$$\text{Last}(x) \wedge \text{LT}(x, y) \rightarrow \text{Halt} \quad (56)$$

$$\text{LT}(x, y) \wedge \text{First}(y) \rightarrow \text{Halt} \quad (57)$$

Lemma 16. \mathcal{R}_6 expresses the query Ω .

Proof. For a database \mathcal{D} over \mathcal{S} , let \mathfrak{M} be the universal model set constructed in Section 3. If $\mathcal{D} \in \Omega$, then $\mathcal{R}_5, \mathcal{D} \models \text{Goal}$ by Theorem 3. Then $\mathcal{R}_6, \mathcal{D} \models \text{Goal}$ since any model of \mathcal{R}_6 and \mathcal{D} must contain Goal by Lemma 15.

Conversely, if $\mathcal{D} \notin \Omega$, then there is $\mathcal{U} \in \mathfrak{M}$ with $\text{Goal} \notin \mathcal{U}$. By Lemma 13, there is a model \mathcal{I} of $\text{brake}(\mathcal{R}_5, \text{Halt})$ with $\mathcal{I}^- = \mathcal{U}$, and hence $\text{Goal} \notin \mathcal{I}$. By construction of \mathfrak{M} , none of the rules (54)–(57) applies to \mathcal{U} , and hence \mathcal{I} is also a model of \mathcal{R}_6 , i.e., $\mathcal{R}_6, \mathcal{D} \not\models \text{Goal}$. \square

Lemma 17. \mathcal{R}_6 is chase-terminating for all databases over the schema \mathcal{S} of the query Ω .

Proof. Consider a chase over \mathcal{R}_6 and input database \mathcal{D} . Chase branches where Halt is eventually part of a node label terminate by Lemma 14. Let b denote any branch of the chase where Halt is not derived, and let \mathcal{I} be the union of all node labels on that branch. We want to show that \mathcal{I} (and hence b) is finite.

By Lemma 13, $\mathcal{I} \models \mathcal{R}_1$. Moreover, since $\text{Halt} \notin \mathcal{I}$, rules (54)–(57) are not applicable to \mathcal{I} . Both properties together suffice to show that the set $\mathcal{I}_{\mathcal{R}_1} = \{p(\vec{n}) \in \mathcal{I} \mid p \text{ is a predicate in } \mathcal{R}_1\}$ is an element of \mathfrak{J}_1 defined in Section 3.

Since the predicates in rule bodies of \mathcal{R}_1 do not occur in any rule head in $\mathcal{R}_5 \setminus \mathcal{R}_1$, we can assume without loss of generality (and without affecting chase termination), that the corresponding rules of $\text{brake}(\mathcal{R}_1, \text{Halt}) \subseteq \mathcal{R}_6$ have been applied first. This shows that \mathcal{I} is equal to the result of a chase with non-disjunctive rules $\mathcal{R}_6 \setminus \text{brake}(\mathcal{R}_1, \text{Halt})$ on a database $\mathcal{I}_{\mathcal{R}_1} \in \mathfrak{J}_1$. The claim follows by noting that any such chase must terminate: this was shown in Section 3, where we described a deterministic process of defining the elements in the universal model set \mathfrak{J}_5 from those in \mathfrak{J}_1 . Each steps in this construction is fully determined and introduces isomorphic sets of nulls irrespectively of the order of rule applications. The only exception are application of rules (15), (16), and (18). For example, given facts $\text{First}(n_1)$, $\text{First}(n_2)$, and $\text{Eq}(n_1, n_2)$, the standard model of Section 3 contains one fact $\text{Root}(u_1)$ with $\text{Rep}(n_1, u_1)$ and $\text{Rep}(n_2, u_1)$, which can be obtained using (15) (on n_1) and (18). If we apply (15) to both n_1 and n_2 before applying (18), we obtain two distinct $\text{Root}(u_1)$ and $\text{Root}(u_1')$. Similar variations can occur with other tree nodes if (16) is applied before (18). It is easy to see that this does not endanger termination, but merely leads to several isomorphic paths in the representation tree. \square

Together, Lemmas 16 and 17 show Theorem 12.

5 Removing Disjunctions

Our main result is that any decidable homomorphism-closed query is expressible by a chase-terminating existential rule set. To conclude the proof of this statement, we remove the disjunction from the rule set \mathcal{R}_6 of Section 4. We present this as a general technique of expressing disjunctive Datalog using existential rules, which is also of independent interest.

For a rule set Σ , the *input schema* $\mathcal{S}_{\text{in}}(\Sigma)$ is the set of all predicates in Σ that do not occur in any rule head. We focus on rule sets that can be split into a disjunctive part and an existential part, such that it is admissible to completely apply the disjunctive rules first, and the existential ones afterwards:

Definition 4. A *split* of a set Σ of disjunctive existential rules consists of a set Σ_1 of disjunctive Datalog rules and a set Σ_2 of existential rules, such that $\Sigma = \Sigma_1 \cup \Sigma_2$ and:

For every database \mathcal{D} over $\mathcal{S}_{\text{in}}(\Sigma)$, and for every chase result \mathfrak{M} over $\langle \Sigma, \mathcal{D} \rangle$, there is a chase result \mathfrak{M}_1 over $\langle \Sigma_1, \mathcal{D} \rangle$, such that $\mathfrak{M} = \{C_{\mathcal{I}} \mid \mathcal{I} \in \mathfrak{M}_1\}$ where each $C_{\mathcal{I}}$ is the (unique) interpretation resulting from some chase over $\langle \Sigma_2, \mathcal{I} \rangle$.

$$\rightarrow \exists w. \text{Init}(w) \wedge \text{Done}(w) \wedge \text{Empty}(w) \quad (58)$$

$$\text{Done}(w) \wedge \text{Init}(w) \wedge \mathbf{p}(\vec{x}) \rightarrow \exists w'. \text{Ins}_{\mathbf{p}}(\vec{x}, w, w') \wedge \text{Subs}(w', w') \wedge \text{Init}(w') \quad (59)$$

$$\text{Done}(w) \wedge \bigwedge_{\mathbf{p}(\vec{x}) \in \beta} \text{Ins}_{\mathbf{p}}(\vec{x}, w, w) \rightarrow \exists w_1. \text{Ins}_{\mathbf{p}_1}(\vec{x}_1, w, w_1) \wedge \text{Subs}(w_1, w_1) \quad (60)$$

$$\text{Done}(w) \wedge \bigwedge_{\mathbf{p}(\vec{x}) \in \beta} \text{Ins}_{\mathbf{p}}(\vec{x}, w, w) \rightarrow \exists w_2. \text{Ins}_{\mathbf{p}_2}(\vec{x}_2, w, w_2) \wedge \text{Subs}(w_2, w_2) \quad (61)$$

$$\text{Ins}_{\mathbf{p}}(\vec{x}, w_0, w_1) \wedge \text{Subs}(w_1, w_2) \rightarrow \text{Ins}_{\mathbf{p}}(\vec{x}, w_2, w_2) \wedge \mathbf{p}'(\vec{x}, w_2) \wedge \text{Subs}(w_0, w_2) \quad (62)$$

$$\text{Empty}(w) \wedge \text{Subs}(w, w') \rightarrow \text{Done}(w') \quad (63)$$

Figure 7: The rule set Σ'_1 , where we instantiate (59) and (62) for all $\mathbf{p} \in \text{Preds}$ in $\Sigma_1 \cup \Sigma_2$, and (60) and (61) for all $\beta \rightarrow \mathbf{p}_1(\vec{x}_1) \vee \mathbf{p}_2(\vec{x}_2) \in \Sigma_1$.

$$\text{Done}(w) \wedge \bigwedge_{\mathbf{p}(\vec{x}) \in \beta} \mathbf{p}'(\vec{x}, w) \rightarrow \exists \vec{z}. \bigwedge_{\mathbf{q}(\vec{y}) \in \eta} \mathbf{q}'(\vec{y}, w) \quad (64)$$

Figure 8: The rule set Σ'_2 , where we instantiate (64) for all $\beta \rightarrow \exists \vec{z}. \eta \in \Sigma_2$.

$$\text{Goal}'(w) \rightarrow \text{Acc}(w) \quad (65)$$

$$\text{Ins}_{\mathbf{p}_1}(\vec{x}_1, w, w_1) \wedge \text{Acc}(w_1) \wedge \text{Ins}_{\mathbf{p}_2}(\vec{x}_2, w, w_2) \wedge \text{Acc}(w_2) \wedge \quad (66)$$

$$\bigwedge_{\mathbf{p}(\vec{x}) \in \beta} \text{Ins}_{\mathbf{p}}(\vec{x}, w, w) \rightarrow \text{Acc}(w) \quad (67)$$

Figure 9: The rule set Σ'_3 , where (66) is instantiated for all rules $\beta \rightarrow \mathbf{p}_1(\vec{x}_1) \vee \mathbf{p}_2(\vec{x}_2) \in \Sigma_1$

Lemma 18. Consider a rule set Σ with split $\langle \Sigma_1, \Sigma_2 \rangle$. There is a set Σ' of existential rules, such that, for every database \mathcal{D} over $\mathcal{S}_{\text{in}}(\Sigma)$, we have:

1. $\mathcal{D}, \Sigma \models \text{Goal}$ iff $\mathcal{D}, \Sigma' \models \text{Goal}$, and
2. if $\langle \Sigma_2, \mathcal{D}_2 \rangle$ is chase-terminating for every database \mathcal{D}_2 over $\mathcal{S}_{\text{in}}(\Sigma_2)$, then $\langle \Sigma', \mathcal{D} \rangle$ is also chase-terminating.

To construct this set Σ' , we assume w.l.o.g. that all disjunctive rules have exactly two disjuncts in the head. We define Σ' as the union of sets Σ'_1 , Σ'_2 and Σ'_3 as shown in Figures 7, 8 and 9, which we explain below.

Σ'_1 uses a technique for modelling sets with chase-terminating existential rules (Krötzsch, Marx, and Rudolph 2019, Fig. 2). We adapt this to sets of ground atoms, called *worlds* and denoted by variables w in the figures. Facts $\text{Ins}_{\mathbf{p}}(\vec{t}, w, w')$ express that world w' is obtained by adding $\mathbf{p}(\vec{t})$ to world w . In particular, $\text{Ins}_{\mathbf{p}}(\vec{t}, w, w)$ states that $\mathbf{p}(\vec{t})$ is in w , and we define $\text{world}(w) = \{\mathbf{p}(\vec{a}) \mid \text{Ins}_{\mathbf{p}}(\vec{a}, w, w) \in \mathcal{I}\}$ for any interpretation \mathcal{I} . Worlds are created by adding database facts (59) or by applying rules to existing worlds (60)–(61). Worlds containing only database facts are marked with *Init*. Predicate *Subs* defines the subset relation on worlds. Rules (62)–(63) copy all prior facts to a new world before marking it *Done*.

Proposition 19. Σ'_1 is chase-terminating and for every \mathcal{D} over $\mathcal{S}_{\text{in}}(\Sigma)$, the (unique) interpretation \mathcal{I} resulting from some chase over Σ'_1 and \mathcal{D} is such that:

- if $\mathbf{p}(\vec{a}) \in \mathcal{D}$ and $\text{Done}(w) \in \mathcal{I}$, there exists w' such that $\text{Ins}_{\mathbf{p}}(\vec{a}, w, w') \in \mathcal{I}$ and $\text{world}(w') = \text{world}(w) \cup \{\mathbf{p}(\vec{a})\}$;

- if $\rho \in \Sigma_1$ is applicable to $\text{world}(w)$, creating $\mathbf{p}_1(\vec{a})$ or $\mathbf{p}_2(\vec{b})$, there exists w_1 and w_2 such that $\{\text{Ins}_{\mathbf{p}_1}(\vec{a}, w, w_1), \text{Ins}_{\mathbf{p}_2}(\vec{b}, w, w_2)\} \subseteq \mathcal{I}$, $\text{world}(w_1) = \text{world}(w) \cup \{\mathbf{p}_1(\vec{a})\}$ and $\text{world}(w_2) = \text{world}(w) \cup \{\mathbf{p}_2(\vec{b})\}$.

Note that we cannot distinguish worlds that are not containing all database facts, and that some worlds may contain more facts than needed to satisfy all disjunctive heads.

Σ'_2 now simulates the application of rules from Σ_2 in any of the worlds. Computations relative to different worlds are independent from each other. Finally, Σ'_3 aggregates results from all worlds: a world is accepting (*Acc*) if either *Goal* was derived locally (65) or it has two successor worlds for a disjunctive rule that are both accepting (66). *Goal* is a consequence if any initial world is accepting (67). This finishes the construction of Σ' as the main ingredient for proving Lemma 18.

Finally, we can apply Lemma 18 to $\mathcal{R}_6 = \text{brake}(\mathcal{R}_5, \text{Halt}) \cup \Pi$ from Section 4, where Π denotes rules (54)–(57). Intuitively, a possible split is $\text{brake}(\mathcal{R}_1, \text{Halt})$ and $\text{brake}(\mathcal{R}_5 \setminus \mathcal{R}_1, \text{Halt}) \cup \Pi$. Formally, however, $\text{brake}(\mathcal{R}_1, \text{Halt})$ is not disjunctive Datalog due to existential rules (2), (3), and (49). However, our result easily extends to such rules with empty body: we can just add them to Σ'_1 and treat their inferences like facts from the initial database. The other properties of Definition 4 are easy to verify. The fact that both rule sets have some common rules, such as (50), is no concern. Finally, it remains to argue termination for Σ_2 as required for item (2) in Lemma 18. This is slightly stronger than Lemma 17 since we must also consider databases that use some inferred predicates of \mathcal{R}_1 . However, the proof of Lemma 17 and the emergency brake technique in general served the main purpose of safeguarding against problematic structures among the inferred predicates of \mathcal{R}_1 , and it is not hard to see that this already showed what we require here. Combining all of our insights, we finally obtain:

Theorem 20. Chase-terminating existential rules capture the class of all decidable homomorphism-closed queries.

6 Limitations of Semi-Decidable Languages

A query language \mathfrak{F} over a schema \mathcal{S} is a function from a set L to $2^{\mathcal{D}_{\mathcal{S}}}$, where $\mathcal{D}_{\mathcal{S}}$ is the set of all databases over schema \mathcal{S} . We say that \mathfrak{F} is *semi-decidable* if membership to L is semi-decidable, and that its *query answering problem* is *decidable* if there exists a TM $M_{\mathfrak{F}}$ that takes as input some

$(l, \mathcal{D}) \in (L \times \mathfrak{D}_S)$ and decides whether $\mathcal{D} \in \mathfrak{F}(l)$.

The set of chase-terminating existential rule sets is a query language that is not semi-decidable (Grahne and Onet 2018) and for which the query answering problem is decidable (by running the chase). In fact, we show that one cannot find a semi-decidable query language with similar properties.

Theorem 21. *There are no semi-decidable query languages that (i) express all decidable, homomorphism-closed queries and (ii) for which query answering is decidable.*

To show this result, we define a set \mathcal{M} of TMs (cf. Definition 5), show that \mathcal{M} can be enumerated up to equivalence if there is a semi-decidable language that satisfies (i) and (ii) above (cf. Lemma 22), and finally prove that the consequence of this implication does not hold (cf. Lemma 23).

Definition 5. Consider the set \mathcal{M} of all TMs M such that: (i) The TM M halts on all inputs. (ii) If M accepts some word w , then w corresponds to a database over schema $\{\text{ed}\}$. (iii) Consider some words w and v that correspond to some \mathcal{D} and \mathcal{E} in $\mathfrak{D}_{\{\text{ed}\}}$, respectively. If M accepts w and there is a homomorphism $h : \mathcal{D} \rightarrow \mathcal{E}$, then M accepts v .

Intuitively, \mathcal{M} is the set of all deciders that solve homomorphism-closed queries over databases in $\mathfrak{D}_{\{\text{ed}\}}$.

Lemma 22. *If there is a semi-decidable query language \mathfrak{F} that satisfies (i) and (ii) in Theorem 21, then \mathcal{M} is enumerable up to equivalence.*

Proof. If there is a language such as \mathfrak{F} , then there is an enumerator P for L that prints out a sequence l_1, l_2, \dots and a decider $M_{\mathfrak{F}}$ that can be used to check if $\mathcal{D} \in \mathfrak{F}(l)$ for each $(l, \mathcal{D}) \in (L \times \mathfrak{D}_{\{\text{ed}\}})$. For each $i \geq 1$, let M_i be the TM that, on input w , performs the following computation: if w corresponds to a database $\mathcal{D} \in \mathfrak{D}_{\{\text{ed}\}}$ and $M_{\mathfrak{F}}$ accepts (l_i, \mathcal{D}) , then *accept*; otherwise, *reject*. By modifying P we can define an enumerator that prints out the sequence M_1, M_2, \dots , which contains \mathcal{M} up to equivalence. \square

Lemma 23. *The set \mathcal{M} is not enumerable up to equivalence.*

Proof Sketch. Assume that there is an enumerator that outputs a sequence M_1, M_2, \dots that includes \mathcal{M} up to equivalence. We obtain a contradiction by defining a sequence $\mathcal{D}_1, \mathcal{D}_2, \dots$ of databases and a TM $M_d \in \mathcal{M}$ that diagonalises over M_1, M_2, \dots and $\mathcal{D}_1, \mathcal{D}_2, \dots$. Namely, for each $i \geq 1$, let $\mathcal{D}_i = \{\text{ed}(u_1, u_2), \dots, \text{ed}(u_{p_{i+1}}, u_1)\}$ where p_{i+1} is the $(i+1)$ -th prime. Moreover, M_d is the TM that, on input w , performs the computation: (1) *Reject* if w does not correspond to some $\mathcal{D} \in \mathfrak{D}_{\{\text{ed}\}}$. (2) *Reject* if \mathcal{D} can be hom-embedded into a path over ed . (3) *Accept* if $\text{ed}(u, u) \in \mathcal{D}$ for some null u . (4) If there is some $i \geq 1$ such that there are less nulls in \mathcal{D}_i than in \mathcal{D} , the TM M_i accepts some serialisation that corresponds to \mathcal{D}_i , and there is a homomorphism $h : \mathcal{D} \rightarrow \mathcal{D}_i$; then *reject*. Otherwise, *accept*. \square

7 Discussion and Conclusion

In this paper, we have established a characterization of all decidable homomorphism-closed Boolean queries. We showed that these are exactly the chase-terminating existential rule queries, that is, queries that can be expressed by a

set of (non-disjunctive) existential rules for which the standard chase universally terminates irrespective of the order of rule applications (as long as it is fair).

By its nature, our result immediately shows that various extensions of our framework do not increase its expressivity:

Theorem 24. *Chase-terminating existential rule queries have the same expressivity as*

1. *existential rule queries with guaranteed existence of some finite chase tree (for every database),*
2. *existential rule queries for which the chase terminates according to some fair strategy (such as datalog-first),*
3. *core-chase-terminating existential rule queries,*
4. *disjunctive chase-terminating existential rule queries.*

Proof. (3) Standard-chase termination implies core-chase termination. On the other hand, core chase termination implies decidability and thus our result applies. (1) and (2) Standard-chase termination implies these weaker form of guarantees, which themselves imply core chase termination. (4) Obviously, every (non-disjunctive) existential rule set is a special case of a disjunctive one and for this special case, disjunctive chase termination coincides with termination of the (non-disjunctive) standard chase. On the other hand, disjunctive existential rule queries are also closed under homomorphisms, and disjunctive universal chase termination obviously implies decidability. So our result applies. \square

However, the applicability of our result does not stop at (syntactic) extensions of our framework, as it applies to arbitrary query languages and querying formalisms of different types. In particular we would like to stress the relationship to the very comprehensive existential rules fragment of *bounded treewidth sets (bts) of rules* (Baget et al. 2011a) that is *not* chase-terminating and encompasses a plethora of well-known existential rule fragments with decidable query entailment, including guarded (Calì, Gottlob, and Kifer 2008), frontier-guarded (Baget et al. 2011a), and glut-guarded existential rules (Krötzsch and Rudolph 2011), as well as greedy bts (Baget et al. 2011b):

Theorem 25. *Let Σ be a bounded-treewidth set of rules and Q a conjunctive query. There is a chase-terminating set Σ_Q of existential rules such that $\mathcal{D}, \Sigma \models Q$ iff $\mathcal{D}, \Sigma_Q \models \text{Goal}$.*

While possibly surprising, this is a straightforward consequence of decidability of conjunctive query entailment from bts and of homomorphism-closedness of existential rule queries in general. Note, however, that every Q would give rise to a different Σ_Q . In fact, asking for a “uniform” chase-terminating existential rules set Σ' satisfying $\mathcal{D}, \Sigma' \models Q$ iff $\mathcal{D}, \Sigma \models Q$ would change the game (Zhang, Zhang, and You 2015). Such a set will not exist in all cases.

While our result addresses many of the open questions regarding *expressivity* of the terminating chase (Krötzsch, Marx, and Rudolph 2019) an important avenue for future work is to investigate potential differences when it comes to the corresponding computational *complexities*. We deem it likely that not all of the discussed chase variants give rise to worst-case optimal computations.

Acknowledgements

This work is partly supported by DFG in project number 389792660 (TRR 248, Center for Perspicuous Systems), by BMBF in the Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI), by the Center for Advancing Electronics Dresden (cfaed), by the ERC Consolidator Grant DeciGUT (project number 771779), and by the ANR project CQFD (ANR-18-CE23-0003).

References

- Abiteboul, S.; Hull, R.; and Vianu, V. 1995. *Foundations of Databases*. Addison-Wesley.
- Baget, J.; Leclère, M.; Mugnier, M.; and Salvat, E. 2011a. On rules with existential variables: Walking the decidability line. *Artif. Intell.* 175(9-10):1620–1654.
- Baget, J.; Mugnier, M.; Rudolph, S.; and Thomazo, M. 2011b. Walking the complexity lines for generalized guarded existential rules. In Walsh (2011), 712–717.
- Beeri, C., and Vardi, M. Y. 1984. A proof procedure for data dependencies. *J. ACM* 31(4):718–741.
- Bourgau, C.; Carral, D.; Krötzsch, M.; Rudolph, S.; and Thomazo, M. 2021. Capturing homomorphism-closed decidable queries with existential rules. arXiv:2107.07811 [cs.LO].
- Cali, A.; Gottlob, G.; and Kifer, M. 2008. Taming the infinite chase: Query answering under expressive relational constraints. In Brewka, G., and Lang, J., eds., *Proc. 11th Int. Conf. on Knowledge Representation and Reasoning (KR'08)*, 70–80. AAAI Press.
- Carral, D.; Dragoste, I.; Krötzsch, M.; and Lewe, C. 2019. Chasing sets: How to use existential rules for expressive reasoning. In Kraus, S., ed., *Proc. 28th Int. Joint Conf. on Artificial Intelligence, IJCAI'19*, 1624–1631. ijcai.org.
- Carral, D.; Dragoste, I.; and Krötzsch, M. 2017. Restricted chase (non)termination for existential rules with disjunctions. In Sierra, C., ed., *Proc. 26th Int. Joint Conf. on Artificial Intelligence, IJCAI'17*, 922–928. ijcai.org.
- Deutsch, A.; Nash, A.; and Rummel, J. B. 2008. The chase revisited. In Lenzerini, M., and Lembo, D., eds., *Proc. 27th Symposium on Principles of Database Systems (PODS'08)*, 149–158. ACM.
- Fagin, R.; Kolaitis, P. G.; Miller, R. J.; and Popa, L. 2005. Data exchange: semantics and query answering. *Theoretical Computer Science* 336(1):89–124.
- Grahne, G., and Onet, A. 2018. Anatomy of the chase. *Fundam. Informaticae* 157(3):221–270.
- Krötzsch, M., and Rudolph, S. 2011. Extending decidable existential rules by joining acyclicity and guardedness. In Walsh (2011), 963–968.
- Krötzsch, M.; Marx, M.; and Rudolph, S. 2019. The power of the terminating chase. In Barceló, P., and Calautti, M., eds., *Proc. 22nd Int. Conf. on Database Theory, ICDT'19*, volume 127 of *LIPIcs*, 3:1–3:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Marnette, B. 2009. Generalized schema-mappings: from termination to tractability. In Paredaens, J., and Su, J., eds., *Proc. 28th Symposium on Principles of Database Systems (PODS'09)*, 13–22. ACM.
- Rudolph, S., and Thomazo, M. 2015. Characterization of the expressivity of existential rule queries. In Yang, Q., and Wooldridge, M., eds., *Proc. 24th Int. Joint Conf. on Artificial Intelligence (IJCAI'15)*, 3193–3199. AAAI Press.
- Walsh, T., ed. 2011. *Proc. 22nd Int. Joint Conf. on Artificial Intelligence (IJCAI'11)*. AAAI Press/IJCAI.
- Zhang, H.; Zhang, Y.; and You, J. 2015. Existential rule languages with finite chase: Complexity and expressiveness. In Bonnet, B., and Koenig, S., eds., *Proc. 29th AAAI Conf. on Artificial Intelligence (AAAI'15)*. AAAI Press.