

Model-theoretic Characterizations of Rule-based Ontologies

Marco Console
Sapienza, University of Rome
Italy
console@diag.uniroma1.it

Phokion G. Kolaitis
UC Santa Cruz & IBM Research
USA
kolaitis@ucsc.edu

Andreas Pieris
University of Edinburgh
UK
apieris@inf.ed.ac.uk

ABSTRACT

An ontology specifies an abstract model of a domain of interest via a formal language that is typically based on logic. Although description logics are popular formalisms for modeling ontologies, tuple-generating dependencies (tgds), originally introduced as a unifying framework for database integrity constraints, and later on used in data exchange and integration, are also well suited for modeling ontologies that are intended for data-intensive tasks. The reason is that, unlike description logics, tgds can easily handle higher-arity relations that naturally occur in relational databases. In recent years, there has been an extensive study of tgd-ontologies and of their applications to several different data-intensive tasks. However, the fundamental question of whether the expressive power of tgd-ontologies can be characterized in terms of model-theoretic properties remains largely unexplored. We establish several characterizations of tgd-ontologies, including characterizations of ontologies specified by such central classes of tgds as full, linear, guarded, and frontier-guarded tgds. Our characterizations use the well-known notions of critical instance and direct product, as well as a novel locality property for tgd-ontologies. We further use this locality property to decide whether an ontology expressed by frontier-guarded (respectively, guarded) tgds can be expressed by tgds in the weaker class of guarded (respectively, linear) tgds, and effectively construct such an equivalent ontology if one exists.

CCS CONCEPTS

• **Theory of computation** → **Logic and databases**; *Description logics*.

KEYWORDS

ontologies; tuple-generating dependencies; guardedness; model theory; finite axiomatizability

ACM Reference Format:

Marco Console, Phokion G. Kolaitis, and Andreas Pieris. 2021. Model-theoretic Characterizations of Rule-based Ontologies. In *Proceedings of the 40th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS '21)*, June 20–25, 2021, Virtual Event, China. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3452021.3458310>



This work is licensed under a Creative Commons Attribution International 4.0 License.

PODS '21, June 20–25, 2021, Virtual Event, China.
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8381-3/21/06.
<https://doi.org/10.1145/3452021.3458310>

1 INTRODUCTION

Model theory is the study of the interaction between formulas in some logical formalism and their models, that is, structures that satisfy the formulas. There are two directions in this interaction, namely, from syntax to semantics and from semantics to syntax. The first direction aims to identify structural properties possessed by all models of formulas having common syntactical features. For example, it is easy to show that every universal first-order sentence is preserved under substructures. The second direction aims to characterize formulas in terms of their structural properties. For example, the well known Łos-Tarski Theorem asserts that if a first-order sentence is preserved under substructures, then it is logically equivalent to a universal first-order sentence. In general, establishing results in the second direction is a much harder task than establishing results in the first. In other words, obtaining model-theoretic characterizations of formulas or of classes of formulas is a far greater challenge than identifying structural properties possessed by all models of formulas with common syntactic features.

Makowsky and Vardi [14] were the first to obtain model-theoretic characterizations of classes of database dependencies expressed in suitable fragments of first-order logic. Furthermore, they classified the work on model-theoretic characterizations into two distinct approaches, which they called the *preservation approach* and the *axiomatizability approach*:

Preservation Approach. In this approach, one considers two logical formalisms \mathcal{L} and \mathcal{L}' , where \mathcal{L}' is typically a proper fragment of \mathcal{L} , and the goal is to obtain model-theoretic characterizations of the following form: a set Σ of \mathcal{L} -formulas is equivalent to a set Σ' of \mathcal{L}' -formulas if and only if the models of Σ satisfy certain structural properties. For example, the aforementioned Łos-Tarski Theorem is a prototypical example in the preservation approach.

Axiomatizability Approach. In this approach, one considers a logical formalism \mathcal{L} , and the goal is to obtain model-theoretic characterizations of the following form: a class C of structures is axiomatizable by a (possibly infinite) set of \mathcal{L} -formulas if and only if the structures in C satisfy certain structural properties. Makowsky and Vardi [14] further distinguished the special case of *finite axiomatizability*, where the goal is to obtain model-theoretic characterizations of when a class of structures is axiomatizable by a finite set of \mathcal{L} -formulas. They then obtained model-theoretic characterizations of axiomatizability and finite axiomatizability of various classes of database dependencies.

As is well known, database dependencies were originally used to formalize integrity constraints in databases with much of the early work in this area focusing on the implication problem between database dependencies (see [10] for a survey). The two most

prominent classes of database dependencies are the class of tuple-generating dependencies (tgds) and the class of equality-generating dependencies (egds). By definition, a tgd is a first-order sentence of the form $\forall \bar{x} \forall \bar{y} (\phi(\bar{x}, \bar{y}) \rightarrow \exists \bar{z} \psi(\bar{x}, \bar{z}))$, where $\phi(\bar{x}, \bar{y})$ is a (possibly empty) conjunction of atoms over a schema S , and $\psi(\bar{x}, \bar{z})$ is a (non-empty) conjunction of atoms over S . Similarly, an egd is a first-order sentence of the form $\forall \bar{x} (\phi(\bar{x}) \rightarrow x_i = x_j)$, where $\phi(\bar{x})$ is a non-empty conjunction of atoms over a schema S and x_i, x_j are variables from \bar{x} . Later on, tgds and egds found numerous uses in other data management tasks. In particular, they have been used as schema-mapping specification languages and, as such, have been (and still are) successfully deployed in the study of data exchange [1, 9] and data integration [11]. More recently, tgds and egds have been used for modeling ontologies.

An ontology specifies an abstract model of a domain of interest using a suitable logical formalism. In particular, description logics have been extensively used as ontology languages [2]. Description logic axioms typically involve unary and binary predicates that correspond, respectively, to concepts in the ontology and roles between two different concepts in the ontology. As it turns out, many axioms used in description logics can be expressed as tgds or egds over relational schemas consisting of unary and binary predicates. Since data-intensive tasks, such as ontology-mediated query answering [4], typically involve predicates of higher arities, it can be argued that tgds and egds are well-suited as specification languages for ontologies. As a matter of fact, tgds have been extensively studied in the context of ontologies under different names, such as Datalog+/- [7] and existential rules [15].

Let us now summarize the known model-theoretic characterizations of tgds, egds, and description logics. Concerning the preservation approach, Lutz et al. [12] obtained model-theoretic characterizations of description logics, i.e., they characterized when a finite set of axioms in certain expressive description logics is equivalent to a finite set of axioms in some less expressive description logic. They also characterized when a first-order sentence is equivalent to a finite set of axioms in a certain description logic. Quite recently, Zhang et al. [17] obtained model-theoretic characterizations of existential rule languages in the preservation approach. In particular, they characterized when a finite set of tgds is equivalent to (i) a finite set of frontier-guarded tgds; (ii) a finite set of guarded tgds; and (iii) a finite set of linear tgds. The preservation approach between two logical formalisms \mathcal{L} and \mathcal{L}' gives rise to a natural decision problem, which we will denote by $\text{Rewrite}(\mathcal{L}, \mathcal{L}')$: given a finite set of Σ of \mathcal{L} -formulas, is there a finite set Σ' of \mathcal{L}' -formulas such that Σ is equivalent to Σ' ? The papers [12] and [17] contain a number of complexity results about the decision problem $\text{Rewrite}(\mathcal{L}, \mathcal{L}')$ for various logical formalisms \mathcal{L} and \mathcal{L}' .

Concerning the axiomatizability and finite axiomatizability approaches, we mentioned earlier that Makowsky and Vardi [14] obtained model-theoretic characterizations of database dependencies. Specifically, the main results in [14] are model-theoretic characterizations of sets of full tgds and egds. Also in the axiomatizability/finite axiomatizability approach, Kolaitis and ten Cate [16] obtained model theoretic characterizations of source-to-target tgds, which are the tgds used to formalize data exchange between a source schema and a target schema. These results notwithstanding, however, the study of model-theoretic characterizations of

sets of arbitrary tgds in the axiomatizability/finite axiomatizability approach has remained largely unexplored so far.

1.1 Summary of Results

Motivated by the preceding state of affairs, we embark here on a systematic investigation of model-theoretic characterizations of classes of tgds in the finite axiomatizability approach. This investigation is carried out in the context of ontologies. From a syntactic point of view, an ontology is specified by a set of formulas in some formalism. From a semantic point of view, an ontology can be identified with the set of all structures (finite or infinite) that satisfy the formulas specifying the ontology. Thus, as a semantic object, an ontology is an isomorphism-closed class of structures (finite or infinite) over some fixed relational schema.

Our goal is to answer the following question: what are necessary and sufficient conditions for an ontology (an isomorphism-closed class of structures) to be specified by a set of tgds? The main outcome of this investigation is to characterize the ontologies that are finitely axiomatizable by a finite set of arbitrary tgds or by a finite set of tgds that belong to one of the main subclasses of tgds, namely, full, frontier-guarded, guarded, and linear tgds.

Our model-theoretic characterizations make use of structural properties encountered in earlier model-theoretic characterizations, such as the ontology being closed under direct products and containing critical structures of every finite cardinality, where a structure is *critical* if each of its relations contains all possible tuples from the domain of the structure. The main innovation, however, is the introduction and use of the notion of (n, m) -locality, where, intuitively, n represents the number of universal quantifiers in the tgds and m represents the number of existential quantifiers in the tgds. Several different notions of locality have been used in earlier model-theoretic characterizations of restricted classes of tgds (e.g., in [14, 16]), yet none of them can be used in characterizations of arbitrary sets of tgds. Our first main result asserts that an ontology O is axiomatizable by a finite set of tgds if and only if O is closed under direct products, contains critical structures of every finite cardinality, and is (n, m) -local for some non-negative integers n and m . The notion of (n, m) -locality turns out to be delicate, yet flexible. Indeed, this notion can be tailored to other classes of tgds, so that it gives rise to the refined notions of frontier-guarded (n, m) -locality, guarded (n, m) -locality, and linear (n, m) -locality. Using these refined notions, we obtain model theoretic characterizations of ontologies that are axiomatizable by a finite set of, respectively, frontier-guarded, guarded, and linear tgds.

Finally, we investigate the decision problems $\text{Rewrite}(\mathcal{L}, \mathcal{L}')$, where \mathcal{L} is the class of frontier-guarded tgds and \mathcal{L}' is the class of guarded tgds, or \mathcal{L} is the class of guarded tgds and \mathcal{L}' is the class of linear tgds. In both these cases, we obtain complexity results that significantly sharpen the results established in [17].

2 PRELIMINARIES

Let C and V be disjoint countably infinite sets of constants and variables, respectively. For $n > 0$, let $[n]$ be the set $\{1, \dots, n\}$.

Relational Instances. A *schema* S is a finite set of relation symbols (or predicates) with associated (positive) arity; we write $\text{ar}(R)$ for

the arity of R . An *instance*¹ I over $S = \{R_1, \dots, R_n\}$, or *S-instance*, is a tuple $(\text{dom}(I), R_1^I, \dots, R_n^I)$, where $\text{dom}(I) \subseteq C$ is a (finite or infinite) domain, and R_1^I, \dots, R_n^I are relations over $\text{dom}(I)$, i.e., $R_i^I \subseteq \text{dom}(I)^{\text{ar}(R_i)}$ for $i \in [n]$. A *fact* of I is an expression of the form $R_i(\bar{c})$, where $\bar{c} \in R_i^I$; let $\text{facts}(I)$ be the set of facts of I . For an S-instance $J = (\text{dom}(J), R_1^J, \dots, R_n^J)$, we write $J \subseteq I$ if $\text{facts}(J) \subseteq \text{facts}(I)$. We say that J is a *subinstance* of I , denoted $J \leq I$, if $\text{dom}(J) \subseteq \text{dom}(I)$, and $R^J = R^I|_{\text{dom}(J)}$ for each $R \in S$ with $R^I|_{\text{dom}(J)}$ being the *restriction of R^I over $\text{dom}(J)$* , i.e., the relation $\{\bar{c} \in R^I \mid \bar{c} \in \text{dom}(J)^{\text{ar}(R)}\}$. Note that $J \leq I$ implies $J \subseteq I$, but the other direction does not hold. A *homomorphism* from I to J is a function $h : \text{dom}(I) \rightarrow \text{dom}(J)$ such that, for each $i \in [n]$, $(c_1, \dots, c_m) \in R_i^I$ implies $(h(c_1), \dots, h(c_m)) \in R_i^J$. We write $h : I \rightarrow J$ for the fact that h is a homomorphism from I to J . We also write $h(\text{facts}(I))$ for the set $\{R(h(\bar{c})) \mid R(\bar{c}) \in \text{facts}(I)\}$. Finally, we say that I and J are *isomorphic*, written $I \simeq J$, if there is an 1-1 homomorphism from I to J such that $h^{-1} : J \rightarrow I$.

Tuple-Generating Dependencies. An atom over S is an expression of the form $R(\bar{v})$, where $R \in S$ and \bar{v} is an $\text{ar}(R)$ -tuple of variables from V . A *tuple-generating dependency* (tgd) σ over a schema S is a constant-free first-order sentence $\forall \bar{x} \forall \bar{y} (\phi(\bar{x}, \bar{y}) \rightarrow \exists \bar{z} \psi(\bar{x}, \bar{z}))$, where $\bar{x}, \bar{y}, \bar{z}$ are tuples of variables of V , $\phi(\bar{x}, \bar{y})$ is a (possibly empty) conjunction of atoms over S , and $\psi(\bar{x}, \bar{z})$ is a (non-empty) conjunction of atoms over S . For brevity, we write σ as $\phi(\bar{x}, \bar{y}) \rightarrow \exists \bar{z} \psi(\bar{x}, \bar{z})$, and use comma instead of \wedge for joining atoms. We refer to $\phi(\bar{x}, \bar{y})$ and $\psi(\bar{x}, \bar{z})$ as the *body* and *head* of σ , denoted $\text{body}(\sigma)$ and $\text{head}(\sigma)$, respectively. By abuse of notation, we may treat a tuple of variables as a set of variables, and a conjunction of atoms as a set of atoms. An instance I satisfies a tgd σ as the one above, written $I \models \sigma$, if the following holds: whenever there exists a function $h : \bar{x} \cup \bar{y} \rightarrow \text{dom}(I)$ such that $h(\phi(\bar{x}, \bar{y})) \subseteq \text{facts}(I)$ (as usual, we write $h(\phi(\bar{x}, \bar{y}))$ for the set $\{R(h(\bar{v})) \mid R(\bar{v}) \in \phi(\bar{x}, \bar{y})\}$), then there exists an extension h' of h such that $h'(\psi(\bar{x}, \bar{z})) \subseteq \text{facts}(I)$. The instance I satisfies a set Σ of tgds, written $I \models \Sigma$, in which case we say that I is a *model* of Σ , if $I \models \sigma$ for each $\sigma \in \Sigma$.

We write $\text{TGD}_{n,m}$ for the family of all possible finite sets of tgds with at most $n \geq 0$ universally, and at most $m \geq 0$ existentially quantified variables.² We also write TGD for the family of all possible finite sets of TGDs, i.e., $\text{TGD} = \bigcup_{n \geq 0, m \geq 0} \text{TGD}_{n,m}$.

Classes of Tuple-Generating Dependencies. In our analysis, we will consider the following central classes of tgds:

Full. A tgd σ is *full* if it has no existentially quantified variables. The class of full tgds, i.e., the family of all possible finite sets of full tgds, is denoted FTGD. Notice that FTGD coincides with the class $\bigcup_{n > 0} \text{TGD}_{n,0}$.

Linear. A tgd σ is *linear* if $\text{body}(\sigma)$ has at most one atom. The class of linear tgds is denoted LTGD.

Guarded. A tgd σ is *guarded* if $\text{body}(\sigma)$ is either empty, or has an atom that contains all the universally quantified variables of σ . The class of guarded tgds is denoted GTGD.

Frontier-Guarded. The *frontier* of a tgd σ , denoted $\text{fr}(\sigma)$, is the set of universally quantified variables occurring in $\text{head}(\sigma)$.

A tgd σ is *frontier-guarded* if $\text{body}(\sigma)$ is either empty, or has an atom that contains all the variables of $\text{fr}(\sigma)$. The class of frontier-guarded tgd is denoted FGTGD.

It is not difficult to verify that

$$\text{LTGD} \subsetneq \text{GTGD} \subsetneq \text{FGTGD}$$

and

$$\text{FGTGD} \neq \text{FTGD}.$$

Given a class C of finite sets of tgds, i.e., $C \subseteq \text{TGD}$, we can naturally define the class $C_{n,m}$, for $n, m \geq 0$, as the class $C \cap \text{TGD}_{n,m}$. For example, $\text{GTGD}_{n,m}$ is the class of guarded tgds with at most $n \geq 0$ universally, and at most $m \geq 0$ existentially quantified variables.

Ontologies. An *ontology* O over a schema S is a (finite or infinite) set of S-instances closed under isomorphisms, i.e., if $I \in O$ and J is an S-instance such that $I \simeq J$, then $J \in O$. Given a class C of tgds, O is a *C-ontology* if there is a set $\Sigma \in C$ such that $I \in O$ iff $I \models \Sigma$.

3 MODEL-THEORETIC PROPERTIES

We proceed to introduce three model-theoretic properties of ontologies that will play a crucial role in our characterizations. In fact, will turn out that those three properties are enough to characterize when an ontology is a TGD-ontology. The first two properties rely on the well-known notions of critical instance and direct product, which have been used in several different contexts (see, e.g., [14]), whereas the third one relies on a novel locality property, which we consider as one of the main conceptual contributions of the present work. In the rest, we fix an arbitrary schema $S = \{R_1, \dots, R_\ell\}$.

3.1 Criticality

An S-instance $I = (\text{dom}(I), R_1^I, \dots, R_\ell^I)$ is called *k-critical*, for some $k > 0$, if $|\text{dom}(I)| = k$, i.e., the domain of I consists of k distinct constants from C , and $R_i^I = \text{dom}(I)^{\text{ar}(R_i)}$ for each $i \in [\ell]$, or, in other words, $\text{facts}(I) = \{R_i(\bar{c}) \mid i \in [\ell] \text{ and } \bar{c} \in \text{dom}(I)^{\text{ar}(R_i)}\}$. For example, the $\{R\}$ -instance I , with R being a binary relation, such that $\text{dom}(I) = \{c, d\}$ and $\text{facts}(I) = \{R(c, c), R(c, d), R(d, c), R(d, d)\}$ is 2-critical. An ontology O over S is *k-critical* if it contains a *k-critical* S-instance. We can now define the notion of critical ontology.

Definition 3.1 (Criticality). An ontology is called *critical* if it is *k-critical* for each integer $k > 0$ ■

It is easy to show that:

LEMMA 3.2. *Every TGD-ontology is critical.*

PROOF. Consider a TGD-ontology O over S . By definition, there exists a finite set Σ of tgds over S such that, for every S-instance I , we have that $I \in O$ iff $I \models \Sigma$. Fix an arbitrary integer $k > 0$, and a *k-critical* S-instance I . We proceed to show that $I \models \Sigma$, which in turn implies that $I \in O$, and thus O is critical. Consider a tgd $\sigma \in \Sigma$ of the form $\phi(\bar{x}, \bar{y}) \rightarrow \exists \bar{z} \psi(\bar{x}, \bar{z})$. It is clear that there exists a function $h : \bar{x} \cup \bar{y} \rightarrow \text{dom}(I)$ such that $h(\phi(\bar{x}, \bar{y})) \subseteq \text{facts}(I)$. Let $c \in \text{dom}(I)$, and consider the extension h' of h such that $h'(z) = c$ for each $z \in \bar{z}$. Since I is *k-critical*, it is straightforward to verify that $h'(\psi(\bar{x}, \bar{z})) \subseteq \text{facts}(I)$, which implies that $I \models \sigma$, as needed. □

¹In mathematical logic literature, instances are called *structures*, which is the term we used in the Introduction. For the remainder of this paper, we adopt the term *instances*.
²A tgd has at least one variable, and thus, $\text{TGD}_{0,0}$ contains only the empty set of tgds.

3.2 Closure Under Direct Products

The *direct product* of the S-instances $I = (\text{dom}(I), R_1^I, \dots, R_\ell^I)$ and $J = (\text{dom}(J), R_1^J, \dots, R_\ell^J)$, denoted $I \otimes J$, is the S-instance $K = (\text{dom}(K), R_1^K, \dots, R_\ell^K)$, where $\text{dom}(K) = \text{dom}(I) \times \text{dom}(J)$, and

$$R_i^K = \left\{ \left((a_1, b_1), \dots, (a_{\text{ar}(R_i)}, b_{\text{ar}(R_i)}) \right) \mid \right. \\ \left. (a_1, \dots, a_{\text{ar}(R_i)}) \in R_i^I \text{ and } (b_1, \dots, b_{\text{ar}(R_i)}) \in R_i^J \right\},$$

for each $i \in [\ell]$. The model-theoretic property of interest follows:

Definition 3.3 (\otimes -closure). An ontology O over S is *closed under direct products* if, for every two S-instances $I, J \in O$, $I \otimes J \in O$. ■

It is not difficult to show the following, which is actually implicit in [8], but we provide a proof for the sake of completeness.

LEMMA 3.4. *Every TGD-ontology is closed under direct products.*

PROOF. Consider a TGD-ontology O over S . By definition, there exists a finite set Σ of tgds over S such that, for every S-instance I , $I \in O$ iff $I \models \Sigma$. Consider two instances $I, J \in O$, and let $K = I \otimes J$. Our goal is to show that $K \in O$, or, equivalently, $K \models \Sigma$. Consider a tgd $\sigma \in \Sigma$ of the form $\phi(\bar{x}, \bar{y}) \rightarrow \exists \bar{z} \psi(\bar{x}, \bar{z})$, and assume that there exists a function $h : \bar{x} \cup \bar{y} \rightarrow \text{dom}(K)$ such that $h(\phi(\bar{x}, \bar{y})) \subseteq \text{facts}(K)$. We show that there exists an extension h' of h such that $h'(\psi(\bar{x}, \bar{z})) \subseteq \text{facts}(K)$. We first observe that there exists a homomorphism h_I from K to I , and a homomorphism h_J from K to J , where, for each $(a, b) \in \text{dom}(K)$, $h_I((a, b)) = a$ and $h_J((a, b)) = b$. This holds since $(a_1, b_1), \dots, (a_n, b_n) \in R^K$ iff $(a_1, \dots, a_n) \in R^I$ and $(b_1, \dots, b_n) \in R^J$, for every $R \in S$. Therefore, $h_I \circ h$ maps $\phi(\bar{x}, \bar{y})$ to $\text{facts}(I)$, and $h_J \circ h$ maps $\phi(\bar{x}, \bar{y})$ to $\text{facts}(J)$. Since $I \models \Sigma$ and $J \models \Sigma$, there is an extension h'_I of $h_I \circ h$, and an extension h'_J of $h_J \circ h$ such that $h'_I(\psi(\bar{x}, \bar{z})) \subseteq \text{facts}(I)$ and $h'_J(\psi(\bar{x}, \bar{z})) \subseteq \text{facts}(J)$. Consider now the extension h' of h such that $h'(z) = (h'_I(z), h'_J(z))$, for every $z \in \bar{z}$. We proceed to show that $h'(\psi(\bar{x}, \bar{z})) \subseteq \text{facts}(K)$. Consider $R(w_1, \dots, w_n) \in \psi(\bar{x}, \bar{z})$, and assume that $h'(R(w_1, \dots, w_n)) = R((a_1, b_1), \dots, (a_n, b_n))$. By the definition of h' , $R(a_1, \dots, a_n) \in \text{facts}(I)$ and $R(b_1, \dots, b_n) \in \text{facts}(J)$. Therefore, by construction, $R((a_1, b_1), \dots, (a_n, b_n)) \in \text{facts}(K)$. □

3.3 Locality

We now proceed to introduce our new locality property of ontologies, which in turn relies on the notion of local embedding of an ontology in an instance. Roughly speaking, an ontology O over S is locally embeddable in an S-instance I if, for every subinstance K of I with a bounded number of active domain elements (i.e., domain elements that occur in $\text{facts}(K)$), we can find an instance $J_K \in O$ such that every local neighbour of K in J_K (i.e., subinstances of J_K that contain K and have a bounded number of additional active domain elements not in $\text{facts}(K)$), can be embedded in I while preserving K . We call the ontology O local if, for every S-instance I , O is locally embeddable in I implies that I belongs to O . We proceed to formalize the above high-level description.

The *active domain* of an instance I , denoted $\text{adom}(I)$, is the set of elements of $\text{dom}(I)$ that occur in at least one fact of I . Consider an S-instance J and a finite set of constants $F \subseteq \text{adom}(J)$. For an integer




Figure 1: O is (n, m) -locally embeddable in I .

$m \geq 0$, the m -neighbourhood of F in J is the set of S-instances

$$\{J' \mid F \subseteq \text{adom}(J'), J' \leq J \text{ and } |\text{adom}(J')| \leq |F| + m\},$$

i.e., all the subinstances of J such that their facts contain constants from F and at most m additional elements not occurring in F . Furthermore, given an S-instance $K \subseteq J$, the m -neighbourhood of K in J is defined as the m -neighbourhood of $\text{adom}(K)$ in J , that is, all the subinstances of J that contain K and their facts mention at most m additional elements not occurring in the facts of K .

Consider an ontology O over S , and an S-instance I . For integers $n, m \geq 0$, we say that O is (n, m) -locally embeddable in I if, for every $K \leq I$ with $|\text{adom}(K)| \leq n$, there is $J_K \in O$ such that $K \subseteq J_K$, and for every J' in the m -neighbourhood of K in J_K , there is a function $h : \text{adom}(J') \rightarrow \text{adom}(I)$, which is the identity on $\text{adom}(K)$, such that $h(\text{facts}(J')) \subseteq \text{facts}(I)$. An illustration of when O is (n, m) -locally embeddable in I is shown in Figure 1; the circles represent the set of facts of the instances. The key property of locality follows.

Definition 3.5 (Locality). An ontology O over S is (n, m) -local, for $n, m \geq 0$, if, for every S-instance I , the following holds: O is (n, m) -locally embeddable in I implies $I \in O$. We further say that O is *local* if there exist $n, m \geq 0$ such that O is (n, m) -local. ■

The next result states that every TGD-ontology is local. Actually, it shows a stronger claim since it relates the integers n, m that witness (n, m) -locality with the number of universally and existentially quantified variables, respectively, that can occur in the tgds.

LEMMA 3.6. *For $n, m \geq 0$, every $\text{TGD}_{n,m}$ -ontology is (n, m) -local.*

PROOF. Consider a $\text{TGD}_{n,m}$ -ontology O over S . By definition, there is a set $\Sigma \in \text{TGD}_{n,m}$ (i.e., each tgd of Σ mentions at most n universally and m existentially quantified variables), such that, for every S-instance I , $I \in O$ iff $I \models \Sigma$. Consider an S-instance I , and assume that O is (n, m) -locally embeddable in I . We proceed to show that $I \in O$, or, equivalently, $I \models \Sigma$. Consider a tgd $\sigma \in \Sigma$ of the form $\phi(\bar{x}, \bar{y}) \rightarrow \exists \bar{z} \psi(\bar{x}, \bar{z})$, and assume that there exists a function $h : \bar{x} \cup \bar{y} \rightarrow \text{dom}(I)$ such that $h(\phi(\bar{x}, \bar{y})) \subseteq \text{facts}(I)$. We show that there exists an extension λ of h such that $\lambda(\psi(\bar{x}, \bar{z})) \subseteq \text{facts}(I)$; the existence of λ is illustrated in Figure 2. Let $K = (\text{dom}(K), R_1^K, \dots, R_\ell^K)$ where $\text{dom}(K)$ is the set of constants occurring in $h(\phi(\bar{x}, \bar{y}))$, and, for each $i \in [\ell]$, $R_i^K = R_i^I|_K$. It is clear that $K \leq I$ with $|\text{adom}(K)| \leq n$ since $\phi(\bar{x}, \bar{y})$ mentions at most n variables. Since, by hypothesis, O is (n, m) -locally embeddable in I , we conclude that there exists $J_K \in O$ such that $K \subseteq J_K$, and, for every J' in the m -neighbourhood of K in J_K , there is a function




Figure 2: The function $\lambda = \mu_L \circ g$ in the proof of Lemma 3.6.

$\mu_{J'} : \text{adom}(J') \rightarrow \text{adom}(I)$, which is the identity on $\text{adom}(K)$, such that $\mu_{J'}(\text{facts}(J')) \subseteq \text{facts}(I)$. It is clear that $h(\phi(\bar{x}, \bar{y})) \subseteq \text{facts}(J_K)$. Since $J_K \in O$, or, equivalently, $J_K \models \Sigma$, there exists an extension g of h such that $g(\psi(\bar{x}, \bar{z})) \subseteq \text{facts}(J_K)$. Let $L = (\text{dom}(L), R_1^L, \dots, R_\ell^L)$ where $\text{dom}(L)$ are the constants occurring in $h(\phi(\bar{x}, \bar{y})) \cup g(\psi(\bar{x}, \bar{z}))$, and, for each $i \in [\ell]$, $R_i^L = R_i^{J_K} \upharpoonright_{\text{dom}(L)}$. It is clear that L is in the m -neighbourhood of K in J_K since \bar{z} has at most m variables. Therefore, there is a function $\mu_L : \text{adom}(L) \rightarrow \text{adom}(I)$, which is the identity on $\text{adom}(K)$, such that $\mu_L(\text{facts}(L)) \subseteq \text{facts}(I)$. Consider the function $\lambda = \mu_L \circ g$. Since g is an extension of h , and μ_L is the identity on the elements occurring in $h(\phi(\bar{x}, \bar{y}))$, we get that $\lambda(v) = h(v)$ for each variable v in $\phi(\bar{x}, \bar{y})$, and thus, λ is an extension of h . Moreover, since $g(\psi(\bar{x}, \bar{z})) \subseteq \text{facts}(L)$, we get that $\lambda(\psi(\bar{x}, \bar{z})) \subseteq \text{facts}(I)$. \square

Domain independence is another central property of ontologies, which will play a crucial role in our technical development, and, as we shall see below, it is guaranteed by locality.

Definition 3.7 (Domain Independence). An ontology O over S is called *domain independent* if, for every two S -instances I and J such that $\text{facts}(I) = \text{facts}(J)$, $I \in O$ iff $J \in O$. \blacksquare

In simple words, O is domain independent if, for every two S -instances that have the same set of facts, but not necessarily the same domain, either are both in O , or none of them is in O . It is not difficult to show that locality implies domain independence.

LEMMA 3.8. *Every local ontology is domain independent.*

PROOF. Consider an ontology O over S that is (n, m) -local for $n, m \geq 0$, and two S -instances I and J such that $\text{facts}(I) = \text{facts}(J)$. We show that $I \in O$ implies $J \in O$; the fact that $J \in O$ implies $I \in O$ is shown analogously. This is done by showing that O is (n, m) -locally embeddable in J , which in turn implies that $J \in O$ since, by hypothesis, O is (n, m) -local. Consider an S -instance $K \leq I$ with $|\text{adom}(K)| \leq n$. Since $\text{facts}(I) = \text{facts}(J)$, it is clear that $K \subseteq I$. Since $I \in O$, it suffices to show that, for every J' in the m -neighbourhood of K in I , there is a function $h_{J'} : \text{adom}(J') \rightarrow \text{adom}(J)$, which is the identity on $\text{adom}(K)$, such that $h_{J'}(\text{facts}(J')) \subseteq \text{facts}(J)$. Since $J' \subseteq I$, which means that $\text{facts}(J') \subseteq \text{facts}(I)$, we get that $\text{facts}(J') \subseteq \text{facts}(J)$. Therefore, $h_{J'}$ is simply the identity on $\text{adom}(J')$. \square

4 CHARACTERIZING TGD-ONTOLOGIES

Are the three main properties presented above, i.e., criticality, ω -closure, and locality, enough to characterize when an ontology is a

TGD-ontology? We proceed to give a positive answer. Actually, we provide a more refined result in the sense that we can characterize when an ontology is a TGD $_{n,m}$ -ontology. Unsurprisingly, (n, m) -locality is the key property for such a refined characterization.

THEOREM 4.1. *Given an ontology O , and integers $n, m \geq 0$, the following are equivalent:*

- (1) O is a TGD $_{n,m}$ -ontology.
- (2) O is critical, closed under direct products, and (n, m) -local.

From the above result, which is interesting in its own right, we obtain a characterization of when an ontology is a TGD-ontology.

COROLLARY 4.2. *Given an ontology O , the following are equivalent:*

- (1) O is a TGD-ontology.
- (2) O is critical, closed under direct products, and local.

The rest of this section is devoted to discussing the proof of Theorem 4.1. Actually, we focus on the non-trivial direction (2) \Rightarrow (1); the direction (1) \Rightarrow (2) follows from Lemmas 3.2, 3.4 and 3.6.

4.1 Some Preparation

We first need to introduce some auxiliary technical notions.

Existential Disjunctive Dependencies. In our proof, we are going to use existential disjunctive dependencies, which essentially generalize tgds with equality and disjunction in the head. More precisely, an *existential disjunctive dependency* (edd) δ over a schema S is a constant-free first-order sentence $\forall \bar{x}(\phi(\bar{x}) \rightarrow \bigvee_{i=1}^k \psi_i(\bar{x}_i))$, where \bar{x} is a tuple of variables of V , $\phi(\bar{x})$ is a (possibly empty) conjunction of atoms over S , and, for each $i \in [k]$, $\bar{x}_i \subseteq \bar{x}$, and $\psi_i(\bar{x}_i)$ is either an equality expression $y = z$ with $\bar{x}_i = \{y, z\}$, or a constant-free formula $\exists \bar{y}_i \chi_i(\bar{x}_i, \bar{y}_i)$ with \bar{y}_i being a tuple of variables from $V \setminus \bar{x}$, and $\chi_i(\bar{x}_i, \bar{y}_i)$ a (non-empty) conjunction of atoms over S . If $k = 1$ and $\psi_1(\bar{x}_1)$ is an equality expression, then δ is called an *equality-generating dependency* (egd). An instance I satisfies the edd δ , written $I \models \delta$, if, whenever there exists a function $h : \bar{x} \rightarrow \text{dom}(I)$ such that $h(\phi(\bar{x})) \subseteq \text{facts}(I)$, then there is $i \in [k]$ such that, if $\psi_i(\bar{y}_i)$ is $y = z$, then $h(y) = h(z)$; otherwise, if $\psi_i(\bar{y}_i)$ is of the form $\exists \bar{y}_i \chi_i(\bar{x}_i, \bar{y}_i)$, then there is an extension h' of h such that $h'(\chi_i(\bar{x}_i, \bar{y}_i)) \subseteq \text{facts}(I)$. The instance I satisfies a set Σ of edds, written $I \models \Sigma$, i.e., I is a *model* of Σ , if $I \models \delta$ for each $\delta \in \Sigma$.

Relative Diagram of an Instance. We now proceed to introduce the diagram of an instance relative to another instance. This can be seen as a refinement of the standard notion of diagram of a relational structure in model theory (see, e.g., [8]). Consider two S -instances I, K such that $K \leq I$, and an integer $\ell \geq 0$. We are interested in the so-called ℓ -diagram of K relative to I , which we define now. Let $A_{K,\ell}$ be the set of all atomic formulas of the form $R(\bar{u})$ that can be formed using predicates from S , constants from $\text{dom}(K)$, and ℓ distinct variables $\star_1, \dots, \star_\ell$ from V , i.e., $R \in S$ and $\bar{u} \in (\text{dom}(K) \cup \{\star_1, \dots, \star_\ell\})^{\text{ar}(R)}$. Let $C_{K,\ell}$ be the set of all (possibly infinite) conjunctions of atomic formulas from $A_{K,\ell}$. Given a formula $\gamma(\bar{\star}) \in C_{K,\ell}$, we can naturally talk about the satisfaction of the sentence $\exists \bar{\star} \gamma(\bar{\star})$ by an instance J , in which case we simply write $J \models \exists \bar{\star} \gamma(\bar{\star})$. The ℓ -diagram of K relative to I , denoted $\Delta_{K,\ell}^I$,

is the (possibly infinite) first-order formula

$$\bigwedge_{\alpha \in \text{facts}(K)} \alpha \wedge \bigwedge_{\substack{c, d \in \text{dom}(K), \\ c \neq d}} \neg(c = d) \wedge \bigwedge_{\substack{\gamma(\bar{\star}) \in C_{K, \ell}, \\ I \not\models \exists \bar{\star} \gamma(\bar{\star})}} \neg \exists \bar{\star} \gamma(\bar{\star}).$$

We are, actually, interested in the first-order formula $\Phi_{K, \ell}^I(\bar{x})$ obtained from $\Delta_{K, \ell}^I$ by replacing each constant element $c \in \text{dom}(K)$ with a new variable $x_c \in \mathbb{V} \setminus \{\star_1, \dots, \star_\ell\}$. As for the formulas of $C_{K, \ell}$, we can naturally talk about the satisfaction of $\exists \bar{x} \Phi_{K, \ell}^I(\bar{x})$ by an instance J , in which case we simply write $J \models \exists \bar{x} \Phi_{K, \ell}^I(\bar{x})$. It is straightforward to verify the following easy lemma:

LEMMA 4.3. *Consider an S-instance I . For each S-instance $K \leq I$ and integer $\ell \geq 0$, it holds that $I \models \exists \bar{x} \Phi_{K, \ell}^I(\bar{x})$.*

4.2 The Proof of (2) \Rightarrow (1)

We now have all the ingredients needed for giving the proof of the direction (2) \Rightarrow (1) in Theorem 4.1. Assume that the ontology O is over a schema \mathbb{S} . The proof proceeds in three main steps:

- (1) We construct a finite set Σ^\vee of edds over \mathbb{S} that mention at most n universally, and at most m existentially quantified variables, in such a way that O is precisely the set of models of Σ^\vee . This exploits the fact that O is 1-critical (since it is critical), and the fact that O is (n, m) -local.
- (2) We then show that there exists a finite set $\Sigma^{\exists=}$ of tgds and egds over \mathbb{S} that is logically equivalent to Σ^\vee , i.e., $\Sigma^{\exists=}$ and Σ^\vee have exactly the same models, written $\Sigma^{\exists=} \equiv \Sigma^\vee$; in fact, $\Sigma^{\exists=}$ is simply the set of tgds and egds in Σ^\vee . This exploits the fact that O is closed under direct products.
- (3) We finally show that there exists a finite set Σ^{\exists} of tgds over \mathbb{S} from $\text{TGD}_{n, m}$ such that $\Sigma^{\exists} \equiv \Sigma^{\exists=}$; in fact, Σ^{\exists} consists of the tgds of $\Sigma^{\exists=}$. This exploits the fact that O is critical.

We proceed to give further details on each of the above steps.

Step 1: The finite set of edds Σ^\vee

Let $E_{n, m}$ be the set that collects all the edds over \mathbb{S} of the form $\forall \bar{x}(\phi(\bar{x}) \rightarrow \bigvee_{i=1}^k \psi_i(\bar{x}_i))$ such that \bar{x} consists of at most n distinct variables, and, for each $i \in [k]$, $\psi_i(\bar{x}_i)$ mentions at most $n + m$ distinct variables. The latter means that, if $\psi_i(\bar{x}_i)$ is a formula of the form $\exists \bar{y}_i \chi_i(\bar{x}_i, \bar{y}_i)$ (i.e., is not an equality expression), then \bar{y}_i consists of at most m distinct variables. It is important to observe that $E_{n, m}$ is finite (up to logical equivalence); this is a consequence of the fact that \mathbb{S} is finite, and the number of variables in each edd of $E_{n, m}$ is finite. We then define the set Σ^\vee of edds as the set of all edds from $E_{n, m}$ that are satisfied by every instance of O , i.e.,

$$\Sigma^\vee = \{\delta \in E_{n, m} \mid \text{for each } I \in O, \text{ it holds that } I \models \delta\}.$$

It is clear that Σ^\vee is finite (up to logical equivalence) since $\Sigma^\vee \subseteq E_{n, m}$. We show that O is precisely the set of models of Σ^\vee .

LEMMA 4.4. *For each S-instance I , $I \in O$ iff $I \models \Sigma^\vee$.*

The (\Rightarrow) direction of Lemma 4.4 holds by construction. We proceed to discuss the non-trivial direction (\Leftarrow) . Consider an S-instance I such that $I \notin O$. By Lemma 3.8, O is domain independent, which allows us to assume that $\text{dom}(I) = \text{dom}(O)$, i.e., all the constants in $\text{dom}(I)$ occur in at least one fact of I . The goal is to show that

$I \not\models \Sigma^\vee$, which in turn establishes Lemma 4.4. We first show the following technical claim that involves $\Phi_{K, m}^I(\bar{x})$; recall that $\Phi_{K, m}^I(\bar{x})$ is the formula obtained from the m -diagram of K relative to I , which here is finite since the domain of K is finite.

CLAIM 4.5. *There exists an S-instance $K \leq I$ with $|\text{adom}(K)| \leq n$ such that, for each $J \in O$, it holds that $J \models \neg \exists \bar{x} \Phi_{K, m}^I(\bar{x})$.*

PROOF. By contradiction, assume that, for each S-instance $K \leq I$ with $|\text{adom}(K)| \leq n$, there exists $J \in O$ such that $J \models \exists \bar{x} \Phi_{K, m}^I(\bar{x})$. We proceed to show that the ontology O is (n, m) -locally embeddable in I , which in turn implies that $I \in O$ since, by hypothesis, O is (n, m) -local. But this contradicts the fact that $I \notin O$.

Fix an arbitrary S-instance $K \leq I$ with $|\text{adom}(K)| \leq n$, and let $J \in O$ be the instance such that $J \models \exists \bar{x} \Phi_{K, m}^I(\bar{x})$. We first observe that $J \models \exists \bar{x} \Phi_{K, m}^I(\bar{x})$ implies the existence of an instance $J_K \subseteq J$ such that $K \simeq J_K$. We can therefore assume, w.l.o.g., that $K \subseteq J$. To show that O is (n, m) -locally embeddable in I , it suffices to show that for every S-instance J' in the m -neighbourhood of K in J , there exists a function $h : \text{adom}(J') \rightarrow \text{adom}(I)$, which is the identity on $\text{adom}(K)$, such that $h(\text{facts}(J')) \subseteq \text{facts}(I)$.

By contradiction, assume that there is J' in the m -neighbourhood of K in J for which there is no function $h : \text{adom}(J') \rightarrow \text{adom}(I)$ that is the identity on $\text{adom}(K)$ with $h(\text{facts}(J')) \subseteq \text{facts}(I)$. Let L be the S-instance defined as the difference between J' and K , i.e., L is such that $\text{facts}(L) = \text{facts}(J') \setminus \text{facts}(K)$, while $\text{dom}(L)$ consists of all the constants occurring in $\text{facts}(J') \setminus \text{facts}(K)$, i.e., $\text{dom}(L) = \text{adom}(L)$. Clearly, there is no function $h : \text{adom}(L) \rightarrow \text{adom}(I)$ that is the identity on $\text{adom}(K)$ such that $h(\text{facts}(L)) \subseteq I$. Observe that $|\text{adom}(L) \setminus \text{adom}(K)| \leq m$; we assume that $\text{adom}(L) \setminus \text{adom}(K) = \{d_1, \dots, d_{m'}\}$ for $m' \leq m$. Let $\gamma(\bar{y})$ be the formula obtained from $\bigwedge_{\alpha \in \text{facts}(L)} \alpha$ after renaming each constant d_i to the variable \star_i ; clearly, $\bar{y} = \star_1, \dots, \star_{m'}$. Since there is no function $h : \text{adom}(L) \rightarrow \text{adom}(I)$ that is the identity on $\text{adom}(K)$ such that $h(\text{facts}(L)) \subseteq I$, we can conclude that $I \not\models \exists \bar{y} \gamma(\bar{y})$. Observe now that, by construction, $\neg \exists \bar{y} \gamma(\bar{y})$ is a conjunct of $\Delta_{K, m}^I$, which in turn implies that the formula $\neg \exists \bar{z} \gamma(\bar{z})$ obtained from $\neg \exists \bar{y} \gamma(\bar{y})$ after renaming each constant $c \in \text{dom}(K) = \text{adom}(K)$ to the variable x_c is a conjunct of $\Phi_{K, m}^I(\bar{x})$. Since $L \subseteq J$, we conclude that $J \models \exists \bar{z} \gamma(\bar{z})$, which in turn implies that $J \not\models \exists \bar{x} \Phi_{K, m}^I(\bar{x})$. But this contradicts the fact that $J \models \exists \bar{x} \Phi_{K, m}^I(\bar{x})$, and the claim follows. \square

Consider now the S-instance K provided by Claim 4.5. We can show that $\neg \exists \bar{x} \Phi_{K, m}^I(\bar{x})$ is logically equivalent to an edd from $E_{n, m}$.

CLAIM 4.6. *There is an edd $\delta \in E_{n, m}$ such that $\delta \equiv \neg \exists \bar{x} \Phi_{K, m}^I(\bar{x})$.*

Actually, the edd δ claimed above is obtained by simply converting $\neg \exists \bar{x} \Phi_{K, m}^I(\bar{x})$ into an equivalent edd via the standard logical transformations. However, we need to argue that, after this transformation, we indeed obtain an edd from $E_{n, m}$, namely (i) the right-hand side is non-empty, (ii) each variable in the right-hand side is either existentially quantified, or appears in the left-hand side, and (iii) there are at most n universally and at most m existentially quantified variables. Item (i) holds since $\Phi_{K, m}^I(\bar{x})$ contains at least one negative conjunct; otherwise, the 1-critical instance in O (it exists since O is 1-critical) is a model of $\exists \bar{x} \Phi_{K, m}^I(\bar{x})$,

which contradicts Claim 4.5. Item (ii) is guaranteed by the fact that $\text{dom}(K) = \text{adom}(K)$ since $\text{dom}(I) = \text{adom}(I)$; recall that the latter relies on the fact that O is domain independent. Finally, item (iii) is ensured by the fact that $|\text{dom}(K)| = |\text{adom}(K)| \leq n$.

Having the above technical claims in place, we can now show that $I \not\models \Sigma^\vee$, which in turn completes the proof of Lemma 4.4. By Claims 4.5 and 4.6, we get that there exists an S -instance $K \leq I$ with $|\text{adom}(K)| \leq n$ such that (i) for every $J \in O$, $J \models \neg \exists \bar{x} \Phi_{K,m}^I(\bar{x})$, and (ii) $\neg \exists \bar{x} \Phi_{K,m}^I(\bar{x})$ is equivalent to an edd $\delta \in E_{n,m}$. By the definition of Σ^\vee , we conclude that $\delta \in \Sigma^\vee$. On the other hand, by Lemma 4.3, $I \models \exists \bar{x} \Phi_{K,m}^I(\bar{x})$, i.e., $I \not\models \delta$, and thus $I \not\models \Sigma^\vee$, as needed.

Step 2: The set of tgds and egds $\Sigma^{\exists,=}$ such that $\Sigma^{\exists,=} \equiv \Sigma^\vee$

We define the set $\Sigma^{\exists,=}$ as the set of all tgds and egds in Σ^\vee , that is,

$$\Sigma^{\exists,=} = \{ \delta \in \Sigma^\vee \mid \delta \text{ is either a tgd or an egd} \}.$$

We show that this is the desired set of tgds and egds

LEMMA 4.7. *It holds that $\Sigma^{\exists,=} \equiv \Sigma^\vee$.*

It is clear that $\Sigma^\vee \models \Sigma^{\exists,=}$, that is, each model of Σ^\vee is a model of $\Sigma^{\exists,=}$, since $\Sigma^{\exists,=}$ is a subset of Σ^\vee . We now discuss the non-trivial direction $\Sigma^{\exists,=} \models \Sigma^\vee$. Towards a contradiction, assume that $\Sigma^{\exists,=} \not\models \Sigma^\vee$. This implies that there is an edd $\delta \in \Sigma^\vee$ of the form $\forall \bar{x}(\phi(\bar{x}) \rightarrow \bigvee_{i=1}^k \psi_i(\bar{x}_i))$ such that $\Sigma^{\exists,=} \not\models \delta$. It is clear that, for each $j \in [k]$, $\sigma_j = \forall \bar{x}(\phi(\bar{x}) \rightarrow \psi_j(\bar{x}_j))$ does not belong to $\Sigma^{\exists,=}$; otherwise, $\Sigma^{\exists,=} \models \delta$ which is not the case. Therefore, by the definition of Σ^\vee , for each $j \in [k]$, there exists an S -instance $I_j \in O$ such that $I_j \not\models \sigma_j$, or, equivalently, $I_j \models \exists \bar{x}(\phi(\bar{x}) \wedge \neg \psi_j(\bar{x}_j))$. We define the S -instance

$$J = I_1 \otimes \cdots \otimes I_k.$$

Since, by hypothesis, O is closed under direct products, we get that $J \in O$. We can also show the following technical claim concerning the instance J .

CLAIM 4.8. *It holds that $J \not\models \delta$.*

By Lemma 4.4, $J \models \Sigma^\vee$, and thus $J \models \delta$ since $\delta \in \Sigma^\vee$. But this contradicts Claim 4.8. Consequently, $\Sigma^{\exists,=} \models \Sigma^\vee$, as needed.

Step 3: The set of tgds Σ^\exists from $\text{TGD}_{n,m}$ such that $\Sigma^\exists \equiv \Sigma^{\exists,=}$

We define the set Σ^\exists as the set of all tgds in $\Sigma^{\exists,=}$, that is,

$$\Sigma^\exists = \{ \delta \in \Sigma^{\exists,=} \mid \delta \text{ is a tgd} \}.$$

We show that this is the desired set of tgds

LEMMA 4.9. *It holds that $\Sigma^\exists \in \text{TGD}_{n,m}$ and $\Sigma^\exists \equiv \Sigma^{\exists,=}$.*

The fact that $\Sigma^\exists \in \text{TGD}_{n,m}$ follows by construction since each edd of Σ^\exists belongs to $E_{n,m}$. It is also easy to see that $\Sigma^{\exists,=} \models \Sigma^\exists$ since $\Sigma^\exists \subseteq \Sigma^{\exists,=}$. For showing that $\Sigma^\exists \models \Sigma^{\exists,=}$, assume by contradiction that there is an egd $\delta \in \Sigma^{\exists,=}$ of the form $\forall \bar{x}(\phi(\bar{x}) \rightarrow y = z)$ such that $\Sigma^\exists \not\models \delta$. This implies that there exists an S -instance I such that $I \not\models \delta$, i.e., there is a function $h : \bar{x} \rightarrow \text{dom}(I)$ such that $h(\phi(\bar{x})) \subseteq \text{facts}(I)$ but $h(y) \neq h(z)$. Observe that $h(\phi(\bar{x})) \not\models \delta$. Let J be a k -critical instance, where k is the number of distinct variables in \bar{x} , such that $h(\phi(\bar{x})) \subseteq \text{facts}(J)$. It is clear that $J \not\models \delta$. Since O is critical, $J \in O$. But this contradicts the fact that $\delta \in \Sigma^{\exists,=} \subseteq \Sigma^\vee$, which means that δ is satisfied by every instance in O .

It should be now clear that the non-trivial direction (2) \Rightarrow (1) of Theorem 4.1 readily follows from Lemmas 4.4, 4.7 and 4.9.

5 CHARACTERIZING FTGD-ONTOLOGIES

The next natural question is whether we can characterize when an ontology can be expressed as a finite set of existential-free tgds, i.e., when an ontology is an FTGD-ontology. Since FTGD is the class $\bigcup_{n>0} \text{TGD}_{n,0}$, an answer to this question can be obtained from the characterization established in the previous section, which exemplifies the usefulness of our new locality property. In particular, the following is an immediate consequence of Theorem 4.1.

COROLLARY 5.1. *Given an ontology O , the following are equivalent:*

- (1) O is an FTGD-ontology.
- (2) O is critical, closed under direct products, and $(n, 0)$ -local for some integer $n > 0$.

Observe also that Theorem 4.1 provides a characterization of when an ontology can be expressed as a finite set of full tgds with at most $n > 0$ universally quantified variables. In particular, given an ontology O and integer $n > 0$, O is an $\text{FTGD}_{n,0}$ -ontology iff O is critical, is closed under direct products, and is $(n, 0)$ -local.

5.1 An Alternative Characterization

At this point, it is worth noting that the question of whether we can provide a characterization for full tgds has been already considered back in the 1980s by Makowsky and Vardi in a slightly different context [14]. As described in the Introduction, the main goal of [14] was to characterize the expressive power of database dependencies in terms of model-theoretic properties. Among others, they studied the question of when a family of databases (i.e., instances with a *finite* domain) can be specified as a (finite or infinite) set of full tgds; note that they did not obtain results about tgds with existentially quantified variables. In principle, the characterization of [14] for full tgds can be lifted to our setting, where we consider finite sets of tgds and unrestricted instances. Unfortunately, the model-theoretic characterization of full tgds in Theorem 3 of [14] turns out to be incorrect since one of the closure properties used, called *closure under duplicating extensions*, does not serve its purpose. In fact, Lemma 7 of [14] states that tgds are closed under duplicating extensions, which, however, is not the case, as we discuss next.

In [14], an instance J is a *duplicating extension* of an instance I if there are constants $c \in \text{dom}(I)$ and $d \notin \text{dom}(I)$ such that

$$\text{dom}(J) = \text{dom}(I) \cup \{d\} \quad \text{and} \quad \text{facts}(J) = \text{facts}(I) \cup h(\text{facts}(I))$$

with $h : \text{dom}(I) \rightarrow \text{dom}(I) \cup \{d\}$ being the identity on $\text{dom}(I) \setminus \{c\}$ and $h(c) = d$.³ Intuitively speaking, J is obtained from I by adding d to $\text{dom}(I)$, and by adding to $\text{facts}(I)$ a copy of itself after renaming c to d . An ontology O over S is *closed under duplicating extensions* if, for every $I \in O$, and S -instance J that is a duplicating extension of I , we have that $J \in O$. Lemma 7 of [14] states that tgds are closed under duplicating extensions, and this has been used to prove the characterization of full tgds in Theorem 3 of [14]. However, we can show that there is an FTGD-ontology *not* closed under duplicating extensions, and thus, Lemma 7 in [14] is incorrect.

³Note that the definition in [14] is for databases, but it can be transferred to instances.

Example 5.2 (Counterexample). Consider the full tgd

$$\sigma = R(x, y), S(y, z) \rightarrow T(x, z).$$

It is clear that the instance I with

$$\text{dom}(I) = \{a, b\} \quad \text{and} \quad \text{facts}(I) = \{R(a, b), S(b, a), T(a, a)\}$$

satisfies σ . It is easy now to verify that the instance J with

$$\begin{aligned} \text{dom}(J) &= \text{dom}(I) \cup \{c\} \quad \text{and} \\ \text{facts}(J) &= \text{facts}(I) \cup \{R(c, b), S(b, c), T(c, c)\} \end{aligned}$$

is a duplicating extension of I due to the function $h : \text{dom}(I) \rightarrow \text{dom}(I) \cup \{c\}$ with $h(a) = c$ and $h(b) = b$. However, $J \not\models \sigma$; indeed, there is a function $g : \{x, y, z\} \rightarrow \text{dom}(J)$ such that $g(\text{body}(\sigma)) = \{R(a, b), S(b, c)\}$, but $g(T(x, z)) = T(a, c) \notin \text{facts}(J)$. ■

In what follows in this section, we obtain an alternative model-theoretic characterization of full tgds that uses a different version of closure under duplicating extensions. The problem with the definition of closure under duplicating extensions in [14] is that it does not distinguish the different occurrences of the special constant, which is mapped to the new constant, appearing in a single fact. Going back to our counterexample, since the different occurrences of the constant a in $T(a, a)$ are not distinguished, there is no way to obtain an atom of the form $T(a, c)$ in the duplicating extension. A valid duplicating extension of I would be the instance J with

$$\begin{aligned} \text{dom}(J) &= \text{dom}(I) \cup \{c\} \quad \text{and} \\ \text{facts}(J) &= \text{facts}(I) \cup \{R(c, b), S(b, c), T(a, c), T(c, a), T(c, c)\}. \end{aligned}$$

This is achieved by the following refined definition. Consider two S -instances I and J . We say that J is a *non-oblivious duplicating extension* of I if there are constants $c \in \text{dom}(I)$ and $d \notin \text{dom}(I)$ such that, for every $R \in S$ and tuple $\bar{t} \in (\text{dom}(I) \cup \{d\})^{\text{ar}(R)}$, $R(\bar{t}) \in J$ iff $h(R(\bar{t})) \in I$ with $h : \text{dom}(I) \cup \{d\} \rightarrow \text{dom}(I)$ being the identity on $\text{dom}(I)$ and $h(d) = c$. The term non-oblivious refers to the fact that now the definition is not oblivious to the different occurrences of the constant c . The desired closure property follows.

Definition 5.3 (Duplicating Extensions Closure). An ontology O over S is *closed under non-oblivious duplicating extensions* if, for every $I \in O$, and S -instance J that is a non-oblivious duplicating extension of I , it holds that $J \in O$. ■

Our alternative characterization of FTGD-ontologies relies on closure under non-oblivious duplicating extensions, 1-criticality, domain independence, and two properties that we have not considered before, namely modularity and closure under intersections. Let us introduce those two properties; let $S = \{R_1, \dots, R_\ell\}$.

Modularity. We start with n -modularity, for $n \geq 0$, which actually provides a small witness instance with at most n domain elements of why an instance does not belong to an ontology. Let us clarify that in [14] there is an analogous notion for classes of databases called n -locality. However, here we adopt the name n -modularity, which has been already used in [16] for a similar notion in the context of schema mappings, in order to avoid any confusion with the new notion of (n, m) -locality introduced in this work. Indeed, (n, m) -locality (even when $m = 0$) is inherently different than the notion of n -modularity, or the notion of n -locality from [14].

Definition 5.4 (Modularity). An ontology O over S is *n -modular*, for some integer $n \geq 0$, if, for every S -instance $I \notin O$, there exists an S -instance $J \leq I$ with $|\text{dom}(J)| \leq n$ such that $J \notin O$. ■

Closure Under Intersections. We now recall closure under intersections. Given two S -instances $I = (\text{dom}(I), R_1^I, \dots, R_\ell^I)$ and $J = (\text{dom}(J), R_1^J, \dots, R_\ell^J)$, the *intersection of I and J* , denoted $I \cap J$, is the instance $(\text{dom}(I) \cap \text{dom}(J), R_1^I \cap R_1^J, \dots, R_\ell^I \cap R_\ell^J)$. Then:

Definition 5.5 (\cap -closure). An ontology O over S is *closed under intersections* if, for every two S -instances $I, J \in O$, $I \cap J \in O$. ■

The Characterization. It is now possible to obtain the alternative characterization of when an ontology is an FTGD-ontology by using the notion of non-oblivious duplicating extension. The proof is similar in spirit to that of Theorem 4.1.

THEOREM 5.6. *Given an ontology O , the following are equivalent:*

- (1) O is an FTGD-ontology.
- (2) O is 1-critical, domain independent, n -modular for some integer $n \geq 0$, closed under intersections, and closed under non-oblivious duplicating extensions.

Let us conclude this section by clarifying that the problematic Theorem 3 in [14] does not use the property of modularity since it follows the axiomatizability approach, where infinite sets of full tgds are allowed. However, here we adopt the finite axiomatizability approach, which means that we are interested only in finite sets of full tgds, and this is the reason why modularity is needed.

6 CHARACTERIZING LTGD-ONTOLOGIES

We now focus our attention on linear tgds, and ask whether we can characterize when an ontology is an LTGD-ontology. In particular, our objective is to obtain characterizations for linear tgds in the spirit of Theorem 4.1 and Corollary 4.2, based on our new locality property. Interestingly, this can be done by replacing locality with a refined version of it, which we call linear locality, whereas criticality and closure under direct products remain untouched.

6.1 Linear Locality

We first refine the notion local embedding by taking into account the fact that linear tgds have at most one body atom. As one might expect, now an ontology O over S is locally embeddable in an S -instance I if, for every $K \subseteq I$ with at most one fact (instead of every subinstance of I) that mentions a bounded number of constants, we can find an instance $J_K \in O$ such that every local neighbour of K in J_K can be embedded in I while preserving K . We proceed to formalize this intuitive description.

Consider an ontology O over a schema S , and an S -instance I . For $n, m \geq 0$, we say that O is *linearly (n, m) -locally embeddable in I* if, for every S -instance $K \subseteq I$ such that $|\text{facts}(K)| \leq 1$ and $|\text{adom}(K)| \leq n$, there exists $J_K \in O$ such that $K \subseteq J_K$, and for every J' in the m -neighbourhood of K in J_K , there is a function $h : \text{adom}(J') \rightarrow \text{adom}(I)$, which is the identity on $\text{adom}(K)$, such that $h(\text{facts}(J')) \subseteq \text{facts}(I)$. The property of linear locality follows.

Definition 6.1 (Linear Locality). An ontology O over S is *linear (n, m) -local*, for $n, m \geq 0$, if, for every S -instance I , O is linearly

(n, m) -locally embeddable in I implies $I \in O$. We also say that O is *linear local* if there are $n, m \geq 0$ such that O is linear (n, m) -local. ■

It is important to observe that linear locality implies locality as this will be crucial for obtaining our main characterization.

LEMMA 6.2. *Consider an ontology O that is linear (n, m) -local for some $n, m \geq 0$. It holds that O is (n, m) -local.*

PROOF. Assume that O is over the schema S . Consider an arbitrary S -instance I such that O is (n, m) -locally embeddable in I . Therefore, by definition, O is linearly (n, m) -locally embeddable in I . Since, by hypothesis, O is linear (n, m) -local, we conclude that $I \in O$, which in turn implies that O is (n, m) -local, as needed. □

6.2 The Characterization

We proceed to show that indeed linear locality is the right notion for obtaining our main characterization. To this end, we first present a technical lemma, dubbed *Linearization Lemma*, that essentially characterizes when a TGD-ontology can be expressed as a finite set of linear tgds.

LEMMA 6.3 (LINEARIZATION). *Given a TGD $_{n,m}$ -ontology O , for some $n, m \geq 0$, the following are equivalent:*

- (1) O is an LTGD-ontology.
- (2) O is an LTGD $_{n,m}$ -ontology.
- (3) O is linear (n, m) -local.

Let us stress that the Linearization Lemma, apart from characterizing when a TGD $_{n,m}$ -ontology is an LTGD $_{n,m}$ -ontology via linear (n, m) -locality, it also tells us the following: if a finite set Σ of tgds can be equivalently rewritten as a finite set Σ_L of linear tgds, then each tgd of Σ_L does not have to use more universally or existentially quantified variables than the tgds of Σ (direction (1) \Rightarrow (2)). This is an interesting fact that, although is not essential for our main characterization concerning linear tgds, it will serve as the basis of the procedure for checking whether a finite set of guarded tgds can be rewritten as a finite set of linear tgds (see Section 9).

Having the Linearization Lemma in place, it is now not difficult to obtain the desired characterizations for linear tgds.

THEOREM 6.4. *Given an ontology O , and integers $n, m \geq 0$, the following are equivalent:*

- (1) O is an LTGD $_{n,m}$ -ontology.
- (2) O is critical, closed under direct prod., and linear (n, m) -local.

PROOF. The fact that (1) \Rightarrow (2) follows from Lemmas 3.2 and 3.4, and the direction (2) \Rightarrow (1) of Lemma 6.3. For (2) \Rightarrow (1), since O is linear (n, m) -local, Lemma 6.2 implies that O is (n, m) -local. Since, by hypothesis, O is also critical and closed under direct products, we get from Theorem 4.1 that O is a TGD $_{n,m}$ -ontology, which in turn allows us to apply the Linearization Lemma (direction (3) \Rightarrow (2)), and get that O is an LTGD $_{n,m}$ -ontology, as needed. □

We conclude this section by observing that Theorem 6.4 provides a characterization of when an ontology is an LTGD-ontology.

COROLLARY 6.5. *Given an ontology O , the following are equivalent:*

- (1) O is an LTGD-ontology.
- (2) O is critical, is closed under direct products, and is linear local.

7 CHARACTERIZING GTGD-ONTOLOGIES

Let us now proceed with guarded tgds, and perform a similar analysis as in the previous section for linear tgds. Our objective here is to obtain characterizations for guarded tgds in the spirit of Theorem 4.1 and Corollary 4.2, by relying on a refined version of our locality property, which we call guarded locality, whereas criticality and closure under direct products remain in place.

7.1 Guarded Locality

We first refine the notion of local embedding by taking into account the fact that guarded tgds have either an empty body, or a body atom that contains all the universally quantified variables. As one might guess, now an ontology O over S is locally embeddable in an S -instance I if, for every guarded subinstance K of I (i.e., K has a fact that contains all the active domain elements), that mentions a bounded number of constants, we can find an instance $J_K \in O$ such that every local neighbour of K in J_K can be embedded in I while preserving the instance K . The formal definition follows.

An instance I is called *guarded* if either $\text{facts}(I) = \emptyset$, or there exists a fact $R(c_1, \dots, c_k) \in \text{facts}(I)$ such that $\text{adom}(I) = \{c_1, \dots, c_k\}$. Consider now an ontology O over a schema S , and an S -instance I . For $n, m \geq 0$, we say that O is *guardedly (n, m) -locally embeddable in I* if, for every guarded S -instance $K \leq I$ with $|\text{adom}(K)| \leq n$, there is $J_K \in O$ such that $K \subseteq J_K$, and for every J' in the m -neighbourhood of K in J_K , there exists a function $h : \text{adom}(J') \rightarrow \text{adom}(I)$, which is the identity on $\text{adom}(K)$, such that $h(\text{facts}(J')) \subseteq \text{facts}(I)$. The property of guarded locality is defined as expected.

Definition 7.1 (Guarded Locality). An ontology O over S is *guarded (n, m) -local*, for $n, m \geq 0$, if, for every S -instance I , the following holds: O is guardedly (n, m) -locally embeddable in I implies $I \in O$. We also say that O is *guarded local* if there are $n, m \geq 0$ such that O is guarded (n, m) -local. ■

As for linear locality, we can show that guarded locality implies locality; the proof is similar to that of Lemma 6.2

LEMMA 7.2. *Consider an ontology O that is guarded (n, m) -local for some $n, m \geq 0$. It holds that O is (n, m) -local.*

7.2 The Characterization

We now show that guarded locality is the right notion for obtaining our main characterization for guarded tgds. To this end, we first present a technical lemma in the spirit of the Linearization Lemma, called *Guardedization Lemma*.

LEMMA 7.3 (GUARDEDIZATION). *Given a TGD $_{n,m}$ -ontology O , for some $n, m \geq 0$, the following are equivalent:*

- (1) O is a GTGD-ontology.
- (2) O is a GTGD $_{n,m}$ -ontology.
- (3) O is guarded (n, m) -local.

Let us note that, similarly to the Linearization Lemma, the Guardedization Lemma, apart from characterizing when a TGD $_{n,m}$ -ontology is a GTGD $_{n,m}$ -ontology via guarded (n, m) -locality, it also tells that, if a finite set Σ of tgds can be equivalently rewritten as a finite set Σ_G of guarded tgds, then each tgd of Σ_G does not have to use more universally or existentially quantified variables than the tgds of Σ ((1) \Rightarrow (2)). This interesting fact will be crucial for the

procedure that checks whether a finite set of frontier-guarded tgds can be rewritten as a finite set of guarded tgds (see Section 9).

Having the Guardedization Lemma in place, it is now easy to obtain the desired characterizations for guarded tgds.

THEOREM 7.4. *Given an ontology O , and integers $n, m \geq 0$, the following are equivalent:*

- (1) O is a GTGD $_{n,m}$ -ontology.
- (2) O is critical, closed under direct prod., and guarded (n, m) -local.

PROOF. The fact that (1) \Rightarrow (2) follows from Lemmas 3.2 and 3.4, and the direction (2) \Rightarrow (3) of Lemma 7.3. For (2) \Rightarrow (1), since O is guarded (n, m) -local, Lemma 7.2 implies that O is (n, m) -local. Since, by hypothesis, O is also critical and closed under direct products, we get from Theorem 4.1 that O is a TGD $_{n,m}$ -ontology, which in turn allows us to apply the Guardedization Lemma (direction (3) \Rightarrow (2)), and get that O is a GTGD $_{n,m}$ -ontology, as needed. \square

We conclude this section by observing that Theorem 7.4 provides a characterization of when an ontology is a GTGD-ontology.

COROLLARY 7.5. *Given an ontology O , the following are equivalent:*

- (1) O is a GTGD-ontology.
- (2) O is critical, is closed under direct prod., and is guarded local.

8 CHARACTERIZING FGTGD-ONTOLOGIES

We now concentrate on frontier-guarded tgds, and provide characterizations similar to those established in the previous sections. This is achieved by exploiting the so-called frontier-guarded locality property (another refinement of our locality property), as well as criticality and closure under direct products.

8.1 Frontier-Guarded Locality

The refined versions of the notion of local embedding for linear and guarded tgds were essentially emerged from the syntactic shape of the tgd bodies. This is somehow also the case for frontier-guarded local embeddings. Frontier-guardedness can be seen as a relativized version of guardedness. Indeed, a frontier-guarded body is essentially a guarded body relative to the frontier, i.e., only the frontier should satisfy the guardedness condition. In the same spirit, one can define frontier-guarded instances, which are essentially guarded instances relative to a certain set of active domain elements.

Consider an instance I , and a finite set $F \subseteq \text{adom}(I)$. We say that I is *guarded relative to F* , or simply *F -guarded*, if either $\text{facts}(I) = \emptyset$, or there is a fact $R(c_1, \dots, c_k) \in \text{facts}(I)$ such that $F \subseteq \{c_1, \dots, c_k\}$. Consider now an ontology O over a schema S , and an S -instance I . For $n, m \geq 0$, we say that O is *fr-guardedly (n, m) -locally embeddable in I* if, for every finite set $F \subseteq \text{adom}(I)$ and F -guarded S -instance $K \leq I$ with $|\text{adom}(K)| \leq n$, there exists $J_K^F \in O$ such that $K \subseteq J_K^F$, and for every J' in the m -neighbourhood of F in J_K^F , there exists a function $h : \text{adom}(J') \rightarrow \text{adom}(I)$, which is the identity on F , such that $h(\text{facts}(J')) \subseteq \text{facts}(I)$. The property of frontier-guarded locality is defined as expected.

Definition 8.1 (Frontier-Guarded Locality). An ontology O over S is *frontier-guarded (n, m) -local*, for $n, m \geq 0$, if, for every S -instance I , O is fr-guardedly (n, m) -locally embeddable in I implies $I \in O$.

We further say that O is *frontier-guarded local* if there are $n, m \geq 0$ such that O is frontier-guarded (n, m) -local. \blacksquare

As for the properties of linear and guarded locality, we can show that frontier-guarded locality implies locality; the proof, which is omitted, is similar to that of Lemma 6.2

LEMMA 8.2. *Consider an ontology O that is frontier-guarded (n, m) -local for some $n, m \geq 0$. It holds that O is (n, m) -local.*

8.2 The Characterization

We now show that frontier-guarded locality is an appropriate notion towards our characterization for frontier-guarded tgds. To this end, we first present a technical lemma in the spirit of the Linearization and Guardedization Lemmas.

LEMMA 8.3. *Given a TGD $_{n,m}$ -ontology O , for some $n, m \geq 0$, the following are equivalent:*

- (1) O is an FGTGD $_{n,m}$ -ontology.
- (2) O is frontier-guarded (n, m) -local.

Having Lemma 8.3 in place, it is now easy to obtain the desired characterizations for frontier-guarded tgds.

THEOREM 8.4. *Given an ontology O , and integers $n, m \geq 0$, the following are equivalent:*

- (1) O is an FGTGD $_{n,m}$ -ontology.
- (2) O is critical, closed under direct products, and frontier-guarded (n, m) -local.

PROOF. The fact that (1) \Rightarrow (2) follows from Lemmas 3.2, 3.4, and 8.3. For (2) \Rightarrow (1), since O is frontier-guarded (n, m) -local, Lemma 8.2 implies that O is (n, m) -local. Since, by hypothesis, O is also critical and closed under direct products, we get from Theorem 4.1 that O is a TGD $_{n,m}$ -ontology, which in turn allows us to apply the Frontier-Guardedization Lemma (direction (2) \Rightarrow (1)), and get that O is an FGTGD $_{n,m}$ -ontology, as needed. \square

We conclude this section by observing that Theorem 8.4 provides a characterization of when an ontology is an FGTGD-ontology.

COROLLARY 8.5. *Given an ontology O , the following are equivalent:*

- (1) O is an FGTGD-ontology.
- (2) O is critical, is closed under direct products, and is frontier-guarded local.

9 RELATIVE EXPRESSIVENESS AND REWRITABILITY

By exploiting our new locality properties, and, in particular, the Linearization and Guardedization Lemmas, we can easily separate, in terms of expressive power, LTGD, GTGD and FGTGD. More precisely, we can devise a finite set of guarded (respectively, frontier-guarded) tgds that cannot be equivalently rewritten as a finite set of linear (respectively, guarded) tgds. Actually, those separations were folklore, and made explicit in the recent work [17]. The value of our analysis is that it provides further insights, using the linear and guarded locality properties, on why those separations hold.

9.1 Semantic Separations

Linear vs. Guarded. Let us first separate LTGD from GTGD. To this end, we need to devise a set $\Sigma_G \in \text{GTGD}$ that is not equivalent to a set $\Sigma_L \in \text{LTGD}$. Consider the set Σ_G consisting of the tgd

$$R(x), P(x) \rightarrow T(x).$$

By the Linearization Lemma (directions (1) \Leftrightarrow (3)), there exists $\Sigma_L \in \text{LTGD}$ such that $\Sigma_G \equiv \Sigma_L$ iff Σ_G is linear (1, 0)-local.⁴ But, we can show that Σ_G is *not* linear (1, 0)-local, and thus, such a set Σ_L does not exist. To this end, we devise the $\{R, P, T\}$ -instance I with

$$\text{dom}(I) = \{c\} \quad \text{and} \quad \text{facts}(I) = \{R(c), P(c)\}.$$

for which it is easy to verify that Σ_G is linearly (1, 0)-locally embeddable in I , but $I \not\models \Sigma_G$. Therefore, Σ_G is not linear (1, 0)-local.

Guarded vs. Frontier-Guarded. Let us now separate GTGD from FGTGD. We need to devise a set $\Sigma_F \in \text{FGTGD}$ that is not equivalent to a set $\Sigma_G \in \text{GTGD}$. Consider the set Σ_F consisting of the tgd

$$R(x), P(y) \rightarrow T(x).$$

By the Guardedization Lemma (directions (1) \Leftrightarrow (3)), there is $\Sigma_G \in \text{GTGD}$ such that $\Sigma_F \equiv \Sigma_G$ iff Σ_F is guarded (2, 0)-local. But, we can show that Σ_F is *not* guarded (2, 0)-local, and thus, such a set Σ_G does not exist. To this end, we devise the $\{R, P, T\}$ -instance I with

$$\text{dom}(I) = \{c\} \quad \text{and} \quad \text{facts}(I) = \{R(c), P(d)\}.$$

for which it is easy to verify that Σ_F is guardedly (2, 0)-locally embeddable in I , but $I \not\models \Sigma_F$. Hence, Σ_F is not guarded (2, 0)-local.

9.2 Rewritability

Having the above semantic separations in place, the next natural question is whether we can decide if a finite set of guarded (respectively, frontier-guarded) tgds can be equivalently rewritten as a finite set of linear (respectively, guarded) tgds. This brings us to the following decision problem; let C_1 and C_2 be classes of tgds:

PROBLEM : Rewrite(C_1, C_2)
 INPUT : A set of tgds $\Sigma \in C_1$.
 QUESTION : Is there a set $\Sigma' \in C_2$ such that $\Sigma \equiv \Sigma'$?

The rest of this section is devoted to studying the problems Rewrite(GTGD, LTGD) and Rewrite(FGTGD, GTGD).

From Guarded to Linear. The problem Rewrite(GTGD, LTGD) has been already considered in the recent work [17]. It was shown to be PSPACE-hard, but no decision procedure was provided. We proceed to pinpoint the complexity of this problem by exploiting the results of this work, in particular, the Linearization Lemma.

THEOREM 9.1. *The following hold:*

- (1) Rewrite(GTGD, LTGD) is 2EXPTIME-complete, and EXPTIME-complete for schemas of bounded arity.
- (2) Given a set $\Sigma \in \text{GTGD}$, a set $\Sigma' \in \text{LTGD}$ such that $\Sigma \equiv \Sigma'$, if one exists, can be computed in double exponential time, and in exponential time in the case of schemas of bounded arity.

The complexity lower bounds claimed in item (1) are shown via a reduction from conjunctive query answering under guarded

⁴By abuse of terminology, we say that Σ_G is linear (1, 0)-local meaning that the ontology consisting of the models of Σ_G is linear (1, 0)-local.

Input: A set $\Sigma \in \text{GTGD}_{n,m}$ for $n, m \geq 0$ over S
Output: A set $\Sigma' \in \text{LTGD}$ such that $\Sigma \equiv \Sigma'$, if one exists; otherwise, \perp

$\Sigma' := \{\sigma \mid \sigma \text{ is over } S, \{\sigma\} \in \text{LTGD}_{n,m} \text{ and } \Sigma \models \sigma\}$
if $\Sigma' \neq \emptyset$ **and** $\Sigma' \models \Sigma$ **then**
 | **return** Σ'
else
 | **return** \perp

Algorithm 1: G-to-L

tgds, which is 2EXPTIME-hard in general, and EXPTIME-hard for schemas of bounded arity [5]. Consider now a set $\Sigma \in \text{GTGD}$ over a schema S , where each tgd of Σ has at most $n \geq 0$ universally and at most $m \geq 0$ existentially quantified variables; it is clear that $\Sigma \in \text{GTGD}_{n,m}$. By the Linearization Lemma, we know that the following statements are equivalent:

- There exists $\Sigma' \in \text{LTGD}$ over S such that $\Sigma \equiv \Sigma'$.
- There exists $\Sigma' \in \text{LTGD}_{n,m}$ over S such that $\Sigma \equiv \Sigma'$.

This essentially means that, even though there are infinitely many finite sets of linear tgds over S , it suffices to search only the fragment of $\text{LTGD}_{n,m}$ over S , which is finite, to find a set Σ' that is equivalent to Σ . This leads to the very simple procedure depicted in Algorithm 1. It first collects in Σ' all the linear tgds over S with at most n universally, and at most m existentially quantified variables, that are entailed by the input set of tgds Σ , and then checks whether Σ' is non-empty and entails Σ ; the latter is actually done by checking whether $\Sigma' \models \sigma$ for each $\sigma \in \Sigma$. We proceed to show that the algorithm G-to-L runs in double exponential time in general, and in single exponential time in the case of schemas of bounded arity, which in turn implies items (1) and (2) of Theorem 9.1.

We first observe that the total number of linear tgds over S with at most n universally quantified variables, and at most m existentially quantified variables, is bounded by

$$\underbrace{|\mathcal{S}| \cdot n^{\text{ar}(\mathcal{S})}}_{\geq \# \text{ of linear bodies}} \cdot \underbrace{2^{|\mathcal{S}| \cdot (n+m)^{\text{ar}(\mathcal{S})}}}_{\geq \# \text{ of heads}}$$

where $\text{ar}(\mathcal{S}) = \max_{R \in \mathcal{S}} \{\text{ar}(R)\}$, while each such linear tgd is of size

$$O\left(\text{ar}(\mathcal{S}) \cdot |\mathcal{S}| \cdot (n+m)^{\text{ar}(\mathcal{S})}\right).$$

We also need to understand the complexity of deciding whether a set of guarded tgds entails a linear tgd (needed in the construction of Σ'), as well as the complexity of deciding whether a set of linear tgds entails a guarded tgd (needed for checking whether $\Sigma' \models \Sigma$). It is easy to see that, given a set of tgds Σ and a single tgd σ of the form $\phi(\bar{x}, \bar{y}) \rightarrow \exists \bar{z} \psi(\bar{x}, \bar{z})$, $\Sigma \models \sigma$ iff Σ and the database D_ϕ , obtain by “freezing” $\phi(\bar{x}, \bar{y})$, entails the Boolean conjunctive query q_ϕ obtained from $\exists \bar{z} \psi(\bar{x}, \bar{z})$ after “freezing” \bar{x} ; see, e.g., [13]. Therefore, the complexity of the implication problems in question can be easily inherited from existing results on conjunctive query answering under guarded and linear tgds [5, 6]:

- Given a set $\Sigma_G \in \text{GTGD}$ and a linear tgd σ_L , both over S , we can decide whether $\Sigma_G \models \sigma_L$ in double exponential time in $\text{ar}(\mathcal{S})$, and single exponential time in the size of Σ_G and σ_L .

<p>Input: A set $\Sigma \in \text{FGTGD}_{n,m}$ for $n, m \geq 0$ over S</p> <p>Output: A set $\Sigma' \in \text{GTGD}$ such that $\Sigma \equiv \Sigma'$, if one exists; otherwise, \perp</p> <p>$\Sigma' := \{\sigma \mid \sigma \text{ is over } S, \{\sigma\} \in \text{GTGD}_{n,m} \text{ and } \Sigma \models \sigma\}$</p> <p>if $\Sigma' \neq \emptyset$ and $\Sigma' \models \Sigma$ then</p> <p> return Σ'</p> <p>else</p> <p> return \perp</p>
--

Algorithm 2: FG-to-G

- Given a set $\Sigma_L \in \text{LTGD}$ and a guarded tgd σ_G , both over S , we can decide whether $\Sigma_L \models \sigma_G$ in exponential time in $\text{ar}(S)$ and the size of σ_G , and in polynomial time in the size of Σ_L .

Putting everything together, we get that the algorithm G-to-L runs in double exponential time in general, and in single exponential time in the case of schemas of bounded arity, as needed.

From Frontier-Guarded to Guarded. Rewrite(FGTGD, GTGD) has been studied in [17], and it was shown to be 2EXPTIME-complete, but without providing an effective procedure that builds an equivalent set of guarded tgds, if one exists. We can provide such a procedure by using the results of this work, in particular, the Guardedization Lemma. We further show that the 2EXPTIME-hardness of Rewrite(FGTGD, GTGD) holds even in the case of bounded arity.

THEOREM 9.2. *The following hold:*

- (1) Rewrite(FGTGD, GTGD) is 2EXPTIME-complete even for schemas of bounded arity.
- (2) Given a set $\Sigma \in \text{FGTGD}$, a set $\Sigma' \in \text{GTGD}$ such that $\Sigma \equiv \Sigma'$, if one exists, can be computed in triple exponential time, and in double exponential time for schemas of bounded arity.

The 2EXPTIME-hardness claimed in item (1), even for schemas of bounded arity, is shown via a reduction from conjunctive query answering under frontier-guarded tgds [3]. Consider now a set $\Sigma \in \text{FGTGD}$ over a schema S , where each tgd of Σ has at most $n \geq 0$ universally and at most $m \geq 0$ existentially quantified variables; it is clear that $\Sigma \in \text{FGTGD}_{n,m}$. By the Guardedization Lemma, the following are equivalent:

- There exists $\Sigma' \in \text{GTGD}$ over S such that $\Sigma \equiv \Sigma'$.
- There exists $\Sigma' \in \text{GTGD}_{n,m}$ over S such that $\Sigma \equiv \Sigma'$.

This leads to the simple procedure depicted in Algorithm 2 that constructs an equivalent set of guarded tgds, if one exists. It first collects in Σ' all the guarded tgds over S with at most n universally, and at most m existentially quantified variables, that are entailed by the input set of tgds Σ , and then checks whether Σ' is non-empty and entails Σ . We proceed to analyze its running time.

We observe that the total number of guarded tgds over S with at most n universally quantified variables, and at most m existentially quantified variables, is bounded by

$$\underbrace{2^{|S| \cdot n^{\text{ar}(S)}}}_{\geq \# \text{ of guarded bodies}} \cdot \underbrace{2^{|S| \cdot (n+m)^{\text{ar}(S)}}}_{\geq \# \text{ of heads}}$$

while each such guarded tgd is of size

$$O\left(\text{ar}(S) \cdot |S| \cdot (n+m)^{\text{ar}(S)}\right).$$

Concerning the implication checks, we can easily inherit from existing results on conjunctive query answering under frontier-guarded and guarded tgds [3, 5] the following complexities:

- Given $\Sigma_F \in \text{FGTGD}$ and a guarded tgd σ_G , both over S , we can decide whether $\Sigma_F \models \sigma_G$ in double exponential time.
- Given $\Sigma_G \in \text{GTGD}$ and a frontier-guarded tgd σ_F , both over S , we can decide whether $\Sigma_G \models \sigma_F$ in double exponential time in $\text{ar}(S)$, and single exponential time in the size of Σ_G .

Putting everything together, we get that FG-to-G runs in triple exponential time in general, and in double exponential time in the case of schemas of bounded arity, and Theorem 9.2 follows.

10 CONCLUDING REMARKS

In this work, we have established model-theoretic characterizations of TGD-ontologies, including characterizations of ontologies specified by central classes of tgds, such as full, linear, guarded, and frontier-guarded tgds. Our characterizations use the well-known properties of criticality and closure under direct products, as well as a novel locality property for TGD-ontologies. We further used this locality property to decide whether an ontology expressed by frontier-guarded (respectively, guarded) tgds can be rewritten as an equivalent one expressed by tgds in the weaker class of guarded (respectively, linear) tgds, and effectively construct such an equivalent ontology if one exists. Although our results (model-theoretic characterizations and rewritability) focus on unrestricted (finite or infinite) instances, it can be shown that they also hold if we concentrate on finite instances. As a next step, we would like to perform a similar analysis that goes beyond TGD-ontologies. In particular, we are planning to consider ontologies specified by tgds, egds, and denial constraints. Moreover, in the context of rewritability, it is interesting to investigate the optimality of the size of the equivalent linear or guarded sets of tgds that we build (whenever they exist).

Acknowledgments. We thank the anonymous referees for their useful feedback. Part of Marco Console's work was done while he was a research associate at the University of Edinburgh. Part of Andreas Pieris' work was done while he was visiting the University of California, Santa Cruz. Marco Console has been partially supported by MIUR under the PRIN 2017 project "HOPE" (prot. 2017MMJJRE), and by the EU under the H2020-EU.2.1.1 project TAILOR, grant id. 952215. Phokion Kolaitis has been partially supported by NSF Award No. 1814152. Andreas Pieris was supported by the EPSRC grant EP/S003800/1 "EQUID".

REFERENCES

- [1] Marcelo Arenas, Pablo Barceló, Leonid Libkin, and Filip Murlak. 2014. *Foundations of Data Exchange*. Cambridge University Press.
- [2] Franz Baader, Ian Horrocks, and Ulrike Sattler. 2009. Description Logics. In *Handbook on Ontologies*. 21–43.
- [3] Jean-François Baget, Marie-Laure Mugnier, Sebastian Rudolph, and Michaël Thomazo. 2011. Walking the Complexity Lines for Generalized Guarded Existential Rules. In *IJCAI*. 712–717.
- [4] Meghyn Bienvenu and Magdalena Ortiz. 2015. Ontology-Mediated Query Answering with Data-Tractable Description Logics. In *Reasoning Web*. 218–307.
- [5] Andrea Cali, Georg Gottlob, and Michael Kifer. 2013. Taming the Infinite Chase: Query Answering under Expressive Relational Constraints. *J. Artif. Intell. Res.* 48 (2013), 115–174.
- [6] Andrea Cali, Georg Gottlob, and Thomas Lukasiewicz. 2012. A general Datalog-based framework for tractable query answering over ontologies. *J. Web Sem.* 14 (2012), 57–83.

- [7] Andrea Cali, Georg Gottlob, Thomas Lukasiewicz, Bruno Marnette, and Andreas Pieris. 2010. Datalog+/-: A Family of Logical Knowledge Representation and Query Languages for New Applications. In *LICS*. 228–242.
- [8] Chen C. Chang and H. Jerome Keisler. 1992. *Model theory, Third Edition*. North-Holland.
- [9] Ronald Fagin, Phokion G. Kolaitis, Renée J. Miller, and Lucian Popa. 2005. Data exchange: semantics and query answering. *Theor. Comput. Sci.* 336, 1 (2005), 89–124.
- [10] Ronald Fagin and Moshe Y. Vardi. 1986. The Theory of Data Dependencies - A Survey. In *Symposia in Applied Mathematics*. 19–71.
- [11] Maurizio Lenzerini. 2002. Data Integration: A Theoretical Perspective. In *PODS*. 233–246.
- [12] Carsten Lutz, Robert Piro, and Frank Wolter. 2011. Description Logic TBoxes: Model-Theoretic Characterizations and Rewritability. In *IJCAI*. 983–988.
- [13] David Maier, Alberto O. Mendelzon, and Yehoshua Sagiv. 1979. Testing Implications of Data Dependencies. *ACM Trans. Database Syst.* 4, 4 (1979), 455–469.
- [14] Johann A. Makowsky and Moshe Y. Vardi. 1986. On the Expressive Power of Data Dependencies. *Acta Inf.* 23, 3 (1986), 231–244.
- [15] Marie-Laure Mugnier and Michaël Thomazo. 2014. An Introduction to Ontology-Based Query Answering with Existential Rules. In *Reasoning Web*. 245–278.
- [16] Balder ten Cate and Phokion G. Kolaitis. 2009. Structural characterizations of schema-mapping languages. In *ICDT*. 63–72.
- [17] Heng Zhang, Yan Zhang, and Guifei Jiang. 2020. Model-theoretic Characterizations of Existential Rule Languages. In *IJCAI*.