

Right-Adjoint for Datalog Programs

ten Cate, Balder; Dalmau, Víctor; Opršal, Jakub

DOI:

[10.4230/LIPIcs.ICDT.2024.10](https://doi.org/10.4230/LIPIcs.ICDT.2024.10)

License:

Creative Commons: Attribution (CC BY)

Document Version

Publisher's PDF, also known as Version of record

Citation for published version (Harvard):

ten Cate, B, Dalmau, V & Opršal, J 2024, Right-Adjoint for Datalog Programs. in G Cormode & M Shekelyan (eds), *27th International Conference on Database Theory (ICDT 2024)*, 10, Leibniz International Proceedings in Informatics, LIPIcs, vol. 290, Schloss Dagstuhl, 27th International Conference on Database Theory, ICDT 2024, Paestum, Italy, 25/03/24. <https://doi.org/10.4230/LIPIcs.ICDT.2024.10>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.


When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Right-Adjoint for Datalog Programs

Balder ten Cate 

Institute for Logic, Language, and Computation, University of Amsterdam, The Netherlands

Víctor Dalmau 

Department of Information and Communication Technologies, Universitat Pompeu Fabra,
Barcelona, Spain

Jakub Opršal 

School of Computer Science, University of Birmingham, UK

Abstract

A Datalog program can be viewed as a syntactic specification of a mapping from database instances over some schema to database instances over another schema. We establish a large class of Datalog programs for which this mapping admits a (generalized) right-adjoint. We employ these results to obtain new insights into the existence of, and methods for constructing, homomorphism dualities within restricted classes of instances. From this, we derive new results regarding the existence of uniquely characterizing data examples for database queries in the presence of integrity constraints.

2012 ACM Subject Classification Theory of computation → Logic and databases

Keywords and phrases Datalog, Adjoints, Homomorphism Dualities, Database Constraints, Conjunctive Queries, Data Examples

Digital Object Identifier 10.4230/LIPIcs.ICDT.2024.10

Funding *Balder ten Cate*: Supported by the European Union’s Horizon 2020 research and innovation programme (MSCA-101031081).

Víctor Dalmau: Supported by the MiCin under grants PID2019-109137GB-C22 and PID2022-138506NB-C22, and the Maria de Maeztu program (CEX2021-001195-M).

Jakub Opršal: Supported by the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No 101034413.

Acknowledgements Thanks to Pierre Bourrhis for detailed feedback and catching mistakes.

1 Introduction

Datalog is a rule-based language for specifying mappings from database instances over an input schema \mathbf{S}_{in} , to database instances over an output schema \mathbf{S}_{out} .

► **Example 1.1.** Consider the Datalog program defined by the following rules:

$$\begin{aligned}\text{Path}(x, y) &:- \text{Edge}(x, y). \\ \text{Path}(x, y) &:- \text{Edge}(x, z), \text{Path}(z, y). \\ \text{Ans}(x, y) &:- \text{Path}(x, y).\end{aligned}$$

This Datalog program takes as input an instance over an input schema $\{\text{Edge}\}$, and produces as output an instance over the schema $\{\text{Ans}\}$, where **Ans** is the transitive closure of **Edge**.

We study the existence of right-adjoints and generalized right-adjoints for Datalog programs, where a *right-adjoint* for a Datalog program P is a function Ω from \mathbf{S}_{out} -instances to \mathbf{S}_{in} -instances, such that for all \mathbf{S}_{in} -instances I and \mathbf{S}_{out} -instances J , $P(I) \rightarrow J$ iff $I \rightarrow \Omega(J)$, where “ \rightarrow ” denotes the existence of a homomorphism. *Generalized* right-adjoints are defined similarly, loosely speaking, except that we allow Ω to map each \mathbf{S}_{out} -instance to a *finite set* of \mathbf{S}_{in} -instances, such that, $P(I) \rightarrow J$ iff $I \rightarrow J'$ for some $J' \in \Omega(J)$. We



© Balder ten Cate, Víctor Dalmau, and Jakub Opršal;
licensed under Creative Commons License CC-BY 4.0

27th International Conference on Database Theory (ICDT 2024).

Editors: Graham Cormode and Michael Shekelyan; Article No. 10; pp. 10:1–10:20

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

identify large classes of Datalog programs for which a right-adjoint, respectively, a generalized right-adjoint, exists. For instance, it will follow from our results that the Datalog program from Example 1.1 has a right-adjoint.

Our motivation for studying (generalized) right-adjoints for Datalog programs comes from the fact that they provide us with a means of constructing homomorphism dualities. A homomorphism duality is a pair (F, D) where F and D are sets of instances, such that an arbitrary instance A admits a homomorphism from an instance in F if and only if A does not admit a homomorphism to any instance in D . In other words, homomorphism dualities equate the existence of a homomorphism of one kind to the non-existence of a homomorphism of another kind. Homomorphism dualities have been studied extensively in the literature on constraint satisfaction problems, and have also found several applications in database theory (e.g., for schema mapping design [2], ontology-mediated data access [6], and query inference from data examples [9, 10]). In particular, in [2, 9], homomorphism dualities were used as a tool for studying the unique characterizability, and exact learnability, of schema mappings and of conjunctive queries. Using the same approach, we can use our results on right-adjoints to derive new results on unique characterizations for (unions of) conjunctive queries in the presence of integrity constraints. In the process, we also obtain a new technique for constructing homomorphism dualities within restricted classes of structures, e.g., transitive digraphs.

► **Contribution 1** (Section 3). We introduce a new fragment of Datalog called *TAM Datalog* (*Tree-Shaped Almost-Monadic Datalog*). We characterize TAM Datalog as a fragment of Monadic Second-Order Logic, and we prove that TAM Datalog is closed under composition.

► **Contribution 2** (Section 4). We show that every connected TAM Datalog program has a right-adjoint, and that every TAM Datalog program has a generalized right-adjoint. We show by means of counterexamples that each of the syntactic conditions imposed by TAM Datalog is necessary for the existence of generalized right-adjoints.

► **Contribution 3** (Section 5). We investigate the relationship between generalized right-adjoints and homomorphism dualities. Generalized right-adjoints can be used for constructing homomorphism dualities. We show that all tree dualities can be accounted for in this way.

► **Contribution 4** (Section 6). Following the approach in [2, 9], we derive new results on unique characterizations for (unions of) conjunctive queries in the presence of integrity constraints. In the process, we obtain a new technique for constructing homomorphism dualities within restricted classes of structures, e.g., transitive digraphs.

Some proofs are omitted and can be found in an appendix of the full version of this paper.

Related Work. Foniok and Tardif [17] studied the existence of right adjoints to Pultr functors which are themselves right adjoints [22] in the special case of digraphs. Translating into our terms a Pultr functor is an interpretation (of digraphs in digraphs) (ϕ_V, ϕ_E) where ϕ_V and ϕ_E are conjunctive queries (with k and $2k$ free variables, respectively, for some $k \geq 1$) defining the output node-set and edge-set respectively. For the special case where ϕ_V just returns the input node-set, it was shown in [17] that the functor defined by (ϕ_V, ϕ_E) has a right adjoint if ϕ_E is connected and acyclic. The setup and characterization were generalized in [13] to arbitrary relational structures. We extensively build on the framework and concepts in [13], but we permit the interpretation to be specified by an arbitrary Datalog program, so that our setup is able to encompass common types of database dependencies.

To our knowledge, this is the first time that adjoints for functors defined by Datalog programs have been studied. Also, it is the first application of functors with right adjoints in the context of unique characterization of database queries. In a different setting, namely approximate graph coloring, the “arc graph” functor was used in [21] where it is additionally argued that functors with a right adjoint, more generally, can play a role in the design and analysis of reductions between (promise) constraint satisfaction problems. The use of Datalog programs for reductions between such problems, although without reference to adjoints, is discussed in [14].

2 Preliminaries

Schemas, Instances, Homomorphisms. A *schema* \mathbf{S} is a finite collection of relation symbols R with specified arity $\text{arity}(R) \geq 0$. An \mathbf{S} -*instance* I is a finite set of facts, where a fact is an expression of the form $R(a_1, \dots, a_n)$ with $R \in \mathbf{S}$ and $n = \text{arity}(R)$. Unless specified otherwise, instances are always assumed to be finite. The *active domain* $\text{adom}(I)$ of I is the set of all values a_i occurring in the facts of I . A *homomorphism* $h : I \rightarrow J$, where I and J are instances over the same schema \mathbf{S} , is a function from $\text{adom}(I)$ to $\text{adom}(J)$ such that the h -image of every fact of I is a fact of J . We will denote by $\text{Inst}[\mathbf{S}]$ the set of all \mathbf{S} -instances.

A k -ary *pointed \mathbf{S} -instance* (for $k \geq 0$) is a pair (I, \mathbf{a}) where I is an \mathbf{S} -instance and \mathbf{a} a k -tuple of elements of $\text{adom}(I)$, called *distinguished elements*. A homomorphism $h : (I, \mathbf{a}) \rightarrow (J, \mathbf{b})$ is a homomorphism $h : I \rightarrow J$ such that $h(\mathbf{a}) = \mathbf{b}$.

Incidence Graph, Connectedness, C-Acyclicity. The *incidence graph* of an instance I is the bipartite multi-graph whose nodes are the elements and the facts of I , and where there is a distinct (undirected) edge (a, f) for every occurrence of the element a in the fact f . We say that an instance is *connected* if its incidence graph is connected, and an instance is *acyclic* if its incidence graph is acyclic. A pointed instance (I, \mathbf{a}) is *c-acyclic* if every cycle in the incidence graph of I contains at least one element from the tuple \mathbf{a} .

Conjunctive Queries and Unions of Conjunctive Queries. For \mathbf{S} a schema and $k \geq 0$, a k -ary *conjunctive query* (CQ) over \mathbf{S} is an expression of the form

$$q(y_1, \dots, y_k) :- \exists \mathbf{x} (\phi_1 \wedge \dots \wedge \phi_n) \quad (\text{Eq. 1})$$

where each ϕ_i is a relational atomic formula, and such that each variable y_i occurs in at least one conjunct ϕ_j . A k -ary *union of conjunctive queries* (UCQ) over \mathbf{S} is a finite disjunction of k -ary CQ s over \mathbf{S} . We denote by $q(I)$ the set of tuples \mathbf{a} for which it holds that $I \models q(\mathbf{a})$.

The *canonical instance* of a CQ of the form (Eq. 1) is the pointed instance (I, \mathbf{y}) where I is the instance with active domain $\{y_1, \dots, y_k, \mathbf{x}\}$ whose facts are the conjuncts of ϕ , and $\mathbf{y} = y_1 \dots y_k$. Conversely, the *canonical CQ* of a pointed instance (I, \mathbf{a}) with $\mathbf{a} = a_1 \dots a_k$, is obtained by associating a unique variable y_a to each $a \in \text{adom}(I)$, letting \mathbf{x} be an enumeration of all variables y_a for $a \in \text{adom}(I) \setminus \{a_1, \dots, a_k\}$, and taking the query $q(y_{a_1}, \dots, y_{a_k}) :- \exists \mathbf{x} \bigwedge_{R(b_1, \dots, b_n) \in I} R(y_{b_1}, \dots, y_{b_n})$. By the well-known Chandra-Merlin theorem, a tuple \mathbf{a} belongs to $q(I)$ if and only if the canonical instance of q homomorphically maps to (I, \mathbf{a}) .

We call a UCQ q *c-acyclic* if the (pointed) canonical instance of each CQ in q is c-acyclic.

Datalog. A Datalog program is specified by a collection of rules, and it defines a mapping from instances over a schema \mathbf{S}_{in} (traditionally known as the EDB schema) to instances over a schema \mathbf{S}_{out} (traditionally known as the IDB schema). The presentation we will give here also allows for auxiliary IDB relations that are not exposed in the output schema.

► **Definition 2.1** (Datalog Program). A Datalog program is a tuple $P = (\mathbf{S}_{in}, \mathbf{S}_{out}, \mathbf{S}_{aux}, \Sigma)$ where $\mathbf{S}_{in}, \mathbf{S}_{out}, \mathbf{S}_{aux}$ are mutually disjoint schemas, and Σ is a set of rules of the form

$$S(\mathbf{x}) :- R_1(\mathbf{y}_1), \dots, R_n(\mathbf{y}_n)$$

where $S \in \mathbf{S}_{out} \cup \mathbf{S}_{aux}$, each $R_i \in \mathbf{S}_{in} \cup \mathbf{S}_{aux}$, and each variable in \mathbf{x} occurs in \mathbf{y}_i for some i .

If P is a Datalog program, then we will often use the notation $\mathbf{S}_{in}^P, \mathbf{S}_{out}^P, \mathbf{S}_{aux}^P$, and Σ^P to refer to the constituents of the tuple P .

The *head* of a rule is the part to the left of the $:-$ sign, and the *body* is the part to the right. The *canonical instance* of a Datalog rule $R_0(\mathbf{x}_0) :- R_1(\mathbf{x}_1), \dots, R_n(\mathbf{x}_n)$ is the pointed instance whose active domain is $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, whose facts are the conjuncts $R_i(\mathbf{x}_i)$ of the rule body, and whose sequence of distinguished elements is the tuple \mathbf{x}_0 . We say that a Datalog program P is *connected* if the canonical instance of each rule is connected. We say that a Datalog program P is *non-recursive* if $\mathbf{S}_{aux}^P = \emptyset$.

If P is a Datalog program and I an \mathbf{S}_{in}^P -instance, then a *solution* for I with respect to P is an instance J over the schema $\mathbf{S}_{in} \cup \mathbf{S}_{out} \cup \mathbf{S}_{aux}$ such that $I \subseteq J$, and such that all the rules of P are satisfied in J (i.e., whenever the body of a rule is satisfied under a variable assignment, then so is the head). The well-known *chase* procedure provides a method for constructing a solution: given a Datalog program P and an \mathbf{S}_{in}^P -instance I , we denote by $\text{chase}_P(I)$ the $\mathbf{S}_{in}^P \cup \mathbf{S}_{out}^P \cup \mathbf{S}_{aux}^P$ -instance obtained from I by applying all rules until convergence. More precisely, $\text{chase}_P(I)$ can be defined as the infinite union $\bigcup_{i \geq 0} \text{chase}_P^i(I)$, where $\text{chase}_P^0(I) = I$, and where $\text{chase}_P^{i+1}(I)$ extends $\text{chase}_P^i(I)$ with all facts that can be derived from facts in $\text{chase}_P^i(I)$ using a rule in Σ^P . We refer to [1] for more details.

► **Lemma 2.2.** For all Datalog programs P and \mathbf{S}_{in}^P -instances I , $\text{chase}_P(I)$ is a solution for I with respect to P . Moreover, it is the intersection of all solutions for I with respect to P .

We denote the \mathbf{S}_{out}^P -reduct of $\text{chase}_P(I)$ by $P(I)$. We say that two Datalog programs P, P' with $\mathbf{S}_{in}^P = \mathbf{S}_{in}^{P'}$ and $\mathbf{S}_{out}^P = \mathbf{S}_{out}^{P'}$ are *equivalent* if, for all \mathbf{S}_{in}^P -instances I , $P(I) = P'(I)$.

By a *Boolean* Datalog program, we mean a Datalog program P where \mathbf{S}_{out}^P consists of a single zero-ary relation symbol, which is customarily denoted as **Ans**. In such cases, write $P(I) = \text{true}$ if $P(I) = \{\text{Ans}()\}$ and $P(I) = \text{false}$ otherwise (i.e., if $P(I) = \emptyset$).

It is well-known that Datalog programs are monotone with respect to homomorphisms:

► **Lemma 2.3.** Let P be any Datalog program, and let I, I' be \mathbf{S}_{in}^P -instances. Every homomorphism $h : I \rightarrow I'$ yields, when restricted to $\text{adom}(P(I))$, a homomorphism from $P(I)$ to $P(I')$.

We can think of the above definition of $P(I)$, in terms of the chase, as a bottom-up account of the semantics of a Datalog program. *Unfoldings* (a.k.a. *expansions*) provide a complementary, top-down account. Given a Datalog program P , the set of *derivable rules* of P is the smallest set of rules that (i) contains all rules of P , and (ii) is closed under the operation of substituting occurrences of rule heads by the corresponding rule bodies (renaming variables as necessary). Given a Datalog program P and a relation $R \in \mathbf{S}_{out}^P$, $\text{Unfoldings}(P, R)$ is the set of canonical instances of derivable rules that have R in the rule head and that only have \mathbf{S}_{in} -relations in the body. Note that this set is in general infinite.

► **Example 2.4.** Let P be the Datalog program consisting of the three rules

$$R(x, y) :- S(x, y) \qquad R(x, x) :- T(x, y) \qquad T(x, y) :- U(x, y), U(y, z)$$

where $\mathbf{S}_{in} = \{U, S\}$, $\mathbf{S}_{out} = \{R\}$, and $\mathbf{S}_{aux} = \{T\}$. Then the derivable rules of P are the rules of P together with the rule $R(x, x) :- U(x, y), U(y, x)$, and $\text{Unfoldings}(P, R)$ consists (up to isomorphism) of the pointed instances $(\{U(a, b), U(b, c)\}, \langle a, a \rangle)$ and $(\{S(a, b)\}, \langle a, b \rangle)$.

► **Lemma 2.5** (Cf. [12]). *For all Datalog programs P , instances $I \in \text{Inst}[\mathbf{S}_{in}^P]$, and \mathbf{S}_{out}^P -facts $R(\mathbf{a})$ over $\text{adom}(I)$, $R(\mathbf{a}) \in P(I)$ iff, for some $(J, \mathbf{b}) \in \text{Unfoldings}(P, R)$, $(J, \mathbf{b}) \rightarrow (I, \mathbf{a})$.*

3 TAM Datalog

TAM Datalog is a fragment of Datalog defined by two requirements: “tree-shaped” and “almost-monadic”. We introduce each in isolation first.

Almost-Monadic Datalog. Recall that a Datalog program is *monadic* if all relations in \mathbf{S}_{aux} are unary. It is well known that monadic Datalog programs can be expressed in Monadic Second-Order logic (MSO). Formally, by a k -ary MSO query over a schema \mathbf{S} , we will mean an MSO formula $\phi(x_1, \dots, x_k)$ over \mathbf{S} . We say that a Datalog program $P = (\mathbf{S}_{in}, \mathbf{S}_{out}, \mathbf{S}_{aux}, \Sigma)$ together with a designated k -ary relation $R \in \mathbf{S}_{out}$, defines an MSO query $\phi_R(\mathbf{x})$ over \mathbf{S}_{in} , if for all \mathbf{S}_{in} -instances I and $\mathbf{a} \in \text{adom}(I)$, $R(\mathbf{a}) \in P(I)$ iff $I \models \phi_R(\mathbf{a})$. The following is folklore in the database literature (cf. [19] for an explicit proof):

► **Theorem 3.1.** *Let P be a monadic Datalog program and $R \in \mathbf{S}_{out}^P$. Then (P, R) defines an MSO query.*

We will now define a weaker restriction, namely that of *almost-monadic* Datalog programs, for which the same holds. These are programs in which every k -ary auxiliary relation has, among its k argument positions, (at most) one specified “*articulation position*”, and the syntax of the rules is constrained in such a way that variables occurring in non-articulation positions can only be used to carry information forward, and not to “perform joins”. Formally:

► **Definition 3.2** (Almost-Monadic Datalog). *An articulation function, for a Datalog program P , is a partial function f mapping relations $R \in \mathbf{S}_{aux}^P$ to a number $f(R) \in \{1, \dots, \text{arity}(R)\}$, which we will call the articulation position of R . Each $i \in \{1, \dots, \text{arity}(R)\}$ other than $f(R)$ is called a non-articulation position of R . A Datalog program is almost-monadic if there exists an articulation function such that, in every rule, each variable occurring in a non-articulation position of an auxiliary relation in a rule body occurs only once in that rule body, and does not occur in the articulation position of an auxiliary relation in the head.*

Note: the articulation conditions pertain to auxiliary relations and not to output relations.

► **Example 3.3.** The Datalog program from Example 1.1 (which outputs all pairs (a, b) for which there is a directed path from a to b) is not monadic but is almost-monadic. The witnessing articulation function assigns to the auxiliary relation **Path** its first position as articulation position. Even if we extend the program with an additional rule **Ans** $(x, y) :- \text{Path}(y, x)$ (so that it computes all pairs (a, b) for which there is a directed path from a to b or from b to a), the resulting program is still almost-monadic. This is because the requirements on the articulation function only pertain to auxiliary relations, not to output relations. On the other hand, if we were to change the rule **Path** $(x, y) :- \text{Edge}(x, z), \text{Path}(z, y)$ to **Path** $(x, y) :- \text{Path}(x, z), \text{Path}(z, y)$, the program would no longer be almost-monadic (cf. also Example 6.2).

For an example of a Datalog program that is *not* almost-monadic, see Example 3.11 below.

► **Proposition 3.4.** *The almost-monadic Datalog program from Example 1.1 is not equivalent to a monadic Datalog program.*

The following result justifies the terminology *almost-monadic*. It shows that almost-monadic Datalog programs can be simulated, in a precise sense, by monadic Datalog programs.

► **Theorem 3.5.** *For each almost-monadic Datalog program P and k -ary relation symbol $R \in \mathbf{S}_{out}^P$, there is a Boolean monadic Datalog program P' where $\mathbf{S}_{in}^{P'} = \mathbf{S}_{in}^P \cup \{Q_1, \dots, Q_k\}$, such that the following are equivalent, for all \mathbf{S}_{in}^P -instances I and $a_1, \dots, a_k \in \text{adom}(I)$:*

1. $R(a_1, \dots, a_k) \in P(I)$,
2. $P'(I \cup \{Q_1(a_1), \dots, Q_k(a_k)\}) = \text{true}$.

► **Example 3.6.** Let P be the Datalog program from Example 1.1. To satisfy the statement of Theorem 3.5, it suffices to define P' as:

$$\begin{aligned} \text{Path}'(x) &:- \text{Edge}(x, y), Q_2(y). \\ \text{Path}'(x) &:- \text{Edge}(x, y), \text{Path}'(y). \\ \text{Ans}() &:- \text{Path}'(x), Q_1(x). \end{aligned}$$

Informally, $\text{Path}'(x)$ holds if there is a path starting at x that ends at a node satisfying Q_2 .

It follows that almost-monadic Datalog is contained in MSO. That is, we have the following analogue of Thm. 3.1 for almost-monadic Datalog programs:

► **Corollary 3.7.** *Let P be an almost-monadic Datalog program and $R \in \mathbf{S}_{out}^P$. Then (P, R) defines an MSO query.*

In summary, we have that almost-monadic Datalog forms a strict extension of monadic Datalog that is still contained in MSO. One may be tempted to conjecture that almost-monadic Datalog is expressively complete for the intersection of Datalog and MSO. However, this is not the case. The easiest way to show this, is using the following Lemma, which is interesting in its own right, as it shows that, when the output schema contains only unary relations, almost-monadic Datalog is no more expressive than monadic Datalog:

► **Lemma 3.8.** *Let P be any almost-monadic Datalog such that every $R \in \mathbf{S}_{out}^P$ is unary. Then P is equivalent to a monadic Datalog program.*

Using this lemma, we can show:

► **Proposition 3.9.** *The unary MSO query “ x lies on a directed R -cycle” is not definable by an almost-monadic Datalog program.*

► **Remark 3.10.** Prop. 3.9 also shows that almost-monadic Datalog is not closed under composition, because the same query can be expressed as the composition of two almost-monadic Datalog programs, where the first computes the transitive closure R^* of the relation R , and the second program consists of the single non-recursive rule $\text{Ans}(x) :- R^*(x, x)$. It also shows that almost-monadic Datalog is strictly included in MODEQ (also known as Flag-and-Check), which is another language contained in the intersection of Datalog and MSO [23]. See also [7] for a characterization of the intersection of MSO and Datalog in terms of infinite domain constraint satisfaction problems. As we will soon see, the intersection of *tree-shaped* Datalog and MSO is (up to logical equivalence) precisely tree-shaped almost-monadic Datalog.

Tree-shapedness. We say that a Datalog program P is *tree-shaped* if the incidence graph of the canonical instance of each rule is acyclic. In particular, since we defined incidence graphs as multigraphs, this implies that no variable occurs twice in the same conjunct in the rule body (but the rule head may contain repeated occurrences of variables). Note that we do not require the incidence graph of the rules to be connected, nor do we make any requirements (say, in the case of binary relations) on the direction of edges. Thus, in tree-shaped Datalog programs, rules such as $T(x) :- R(x, y), S(x, y)$ or $T(x) :- R(x, x)$ are forbidden.

► **Example 3.11.** Consider the tree-shaped Datalog program P given by the following two rules (where \mathbf{S}_{in}^P consists of two binary relations, E, F):

$$R(x, y) :- E(x, u), F(u, y) \quad R(x, y) :- E(x, u), R(u, v), F(v, y) \quad \mathbf{Ans}(x, y) :- R(x, y)$$

Then $\mathbf{Ans}^{P(I)}$ contains all pairs (a, b) , such that there is a directed path from a to b in I consisting of a number of E -edges followed by an equal number of F -edges.

Observe that P is not TAM Datalog because neither the first position of the relation R qualifies as an articulation position (since v occurs twice in the second rule body) nor the second position (since u occurs twice in the same rule body).

It follows from known facts about MSO (viz. the fact that MSO on words captures the regular languages) that (P, \mathbf{Ans}) does not define an MSO query. In particular, P is not equivalent to a monadic Datalog program, or even an almost-monadic Datalog program.

► **Lemma 3.12.** *Let P be any tree-shaped Datalog program. Then, for each $R \in \mathbf{S}_{out}^P$, $\text{Unfoldings}(P, R)$ consists of acyclic pointed instances.*

TAM Datalog. A *TAM Datalog program* is a tree-shaped, almost-monadic Datalog program. We will give a precise model-theoretic characterization of TAM Datalog in terms of MSO.

We say that an MSO query $\phi(\mathbf{x})$ is *tree-determined* if for each pointed instance (I, \mathbf{a}) , we have that $I \models \phi(\mathbf{a})$ if and only if there is an acyclic pointed instance (J, \mathbf{b}) such that $J \models \phi(\mathbf{b})$ and $(J, \mathbf{b}) \rightarrow (I, \mathbf{a})$. Note that J must be finite and that J is not required to be connected.

► **Theorem 3.13.** *Let $\phi(x_1, \dots, x_n)$ be an MSO formula. The following are equivalent:*

1. ϕ is definable by a TAM Datalog program,
2. ϕ is definable by a tree-shaped Datalog program,
3. ϕ is tree-determined.

► **Remark 3.14.** It is worth comparing this to the result in [19] that states that monadic Datalog and MSO have the same expressive power on finite trees. Besides the fact that Thm. 3.13 is a characterization on arbitrary (finite) instances while the result in [19] is restricted to trees, there are a few other important differences: in [19], it is assumed that trees are represented as structures in which the children of each node are ordered; that the signature includes predicates marking the root, leafs, the first child of each node, and the last child of each node; and that each node of the tree is labeled by precisely one of the (other) unary predicates in the signature. These assumptions together imply that every homomorphism between such trees is necessarily an isomorphism, which makes the two results incomparable.

► **Corollary 3.15** (TAM Datalog is closed under composition). *For all TAM Datalog programs P_1 and P_2 with $\mathbf{S}_{in}^{P_2} = \mathbf{S}_{out}^{P_1}$, there is a TAM Datalog program $P_3 = (\mathbf{S}_{in}^{P_1}, \mathbf{S}_{out}^{P_2}, \mathbf{S}'_{aux}, \Sigma')$ such that, for all $\mathbf{S}_{in}^{P_1}$ -instances I , $P_3(I) = P_2(P_1(I))$.*

We also provide a syntactic normal form for TAM Datalog programs. A TAM Datalog program is *simple* if every rule body contains precisely one occurrence of a relation from \mathbf{S}_{in} . For instance the program given in Example 1.1 is a simple TAM Datalog program.

► **Theorem 3.16.** *Every (connected) TAM Datalog program can be transformed in polynomial-time into an equivalent (connected) simple TAM Datalog program.*

We make use of this normal form in some of our proofs.

4 Right-Adjoint for TAM Datalog

The notion of adjunction comes from category theory. Although for the most part, we do not assume that the reader has a background in category theory, in order to motivate our definition of generalized right-adjoints for Datalog programs, it is helpful to briefly discuss Datalog programs from a categorical perspective.

Recall that each Datalog program P defines a mapping from $\text{Inst}[\mathbf{S}_{in}^P]$ to $\text{Inst}[\mathbf{S}_{out}^P]$, where $\text{Inst}[\mathbf{S}]$ denotes the set of all \mathbf{S} -instances. Recall also that this mapping is monotone with respect to homomorphisms (cf. Lemma 2.3). We view $\text{Inst}[\mathbf{S}_{in}^P]$ and $\text{Inst}[\mathbf{S}_{out}^P]$ as partial (pre)orders, which can consequently be viewed as *thin categories* where the objects are the \mathbf{S} -instances and there is an arrow from I to J if there exists a homomorphism $h : I \rightarrow J$. The categorical notion of a *functor* is then simply a monotone mapping. In particular, each Datalog program P defines a functor. For functors $F : X \rightarrow Y$ and $G : Y \rightarrow X$, where X and Y are arbitrary thin categories, it is said that G is a *right-adjoint* for F , and that F is a *left-adjoint* of G , if it holds that $F(I) \rightarrow J$ iff $I \rightarrow G(J)$.¹

In this section, we study the existence of right-adjoints for Datalog programs.

► **Example 4.1.** Consider the Datalog program $P = (\mathbf{S}_{in}, \mathbf{S}_{out}, \emptyset, \Sigma)$, where $\mathbf{S}_{in} = \{R\}$, $\mathbf{S}_{out} = \{S\}$, and Σ consists of the rules $S(x, y) :- R(x, y)$ and $S(x, y) :- R(y, x)$. If we think of an input instance I as a directed graph, $P(I)$ is its *symmetric closure*. For every \mathbf{S}_{out}^P -instance J , let $\Omega(J)$ be the \mathbf{S}_{in}^P -instance that is the *maximal symmetric sub-instance* of J , that is, $\Omega(J)$ consists of all facts $R(x, y)$ for which it holds that J contains both $S(x, y)$ and $S(y, x)$. It is not hard to see that $P(I) \rightarrow J$ iff $I \rightarrow \Omega(J)$. Hence, Ω is a right-adjoint of P .

► **Example 4.2.** Consider the TAM Datalog program $P = (\mathbf{S}_{in}, \mathbf{S}_{out}, \emptyset, \Sigma)$, where $\mathbf{S}_{in} = \{Q_1, Q_2\}$, $\mathbf{S}_{out} = \{Q_3\}$, and Σ consists of the rule $Q_3() :- Q_1(x), Q_2(y)$. This Datalog program does *not* have a right-adjoint in the above sense. Indeed, let J be the empty instance. Then $P(I) \rightarrow J$ holds if and only if either I has no Q_1 -facts or I has no Q_2 -facts, a condition that cannot be equivalently characterized by the existence of a homomorphism from I to any fixed single instance J' . However, it can be shown that $P(I) \rightarrow J$ if and only if either $I \rightarrow J'_1$ or $I \rightarrow J'_2$, where $J'_1 = \{Q_1(a)\}$ and $J'_2 = \{Q_2(a)\}$. If we generalize the notion of right-adjoint by allowing $\Omega(J)$ to be a finite set of instances, then, as we will see later (Thm. 4.5), P *does* admit such a right-adjoint, and the fact that $\Omega(J)$ needs to consist of multiple instances is related to the fact that the program is not connected.

Motivated by the above examples and other considerations that will become clear soon, the precise notion of right-adjoints that we will adopt here is a little more refined:

¹ Note that this coincides with the usual definition of adjoint functors as long as both categories are thin. It is also precisely the notion of right-adjoints for Galois connections over preordered sets. See also Remark 4.11 for why we adopt this “thin” definition of adjoints.

► **Definition 4.3** (Generalized Right-Adjoint). A generalized right-adjoint for a Datalog program P is a function Ω_P that maps each $J \in \text{Inst}[\mathbf{S}_{out}^P]$ to a finite set of pairs (J', ι) where $J' \in \text{Inst}[\mathbf{S}_{in}^P]$ and $\iota : J' \rightarrow J$ is a partial function, such that the following holds: for all $I \in \text{Inst}[\mathbf{S}_{in}^P]$, there is a homomorphism $h : P(I) \rightarrow J$ iff there is a homomorphism $h' : I \rightarrow J'$ for some $(J', \iota) \in \Omega_P(J)$, and, furthermore, the homomorphism h' , respectively h , can be chosen such that the following diagram commutes.²

$$\begin{array}{ccc} \text{adom}(P(I)) & \xrightarrow{h} & \text{adom}(J) \\ id \uparrow & & \uparrow \iota \\ \text{adom}(I) & \xrightarrow{h'} & \text{adom}(J') \end{array}$$

Here, the “ \rightarrow ” arrow refers to homomorphisms, and “ \rightarrow ” is used to refer to a partial function. The partial functions ι are needed later on to reason about pointed instances (cf. for instance the proof of Theorem 5.4).

This notion of generalized right-adjoint behaves as one would expect. In particular, if two Datalog programs have generalized right-adjoints, then so does their composition.

► **Example 4.4.** Let P be the Datalog program with input schema $\{R_1, R_2\}$ and output schema $\{Q_1, Q_2\}$ and consisting of the rules

$$Q_1(x) :- R_1(x) \qquad Q_1(x) :- R_2(x) \qquad Q_2(x) :- R_1(x), R_2(x)$$

This Datalog program indeed has a generalized right-adjoint Ω_P . In fact it has a regular right-adjoint, in the sense that $\Omega_P(J)$ is a singleton for all $J \in \text{Inst}[\mathbf{S}_{out}^P]$, as will follow from Theorem 4.5 below. We will illustrate this with an example. Let J be the \mathbf{S}_{out}^P -instance consisting of the single fact $Q_1(a)$. Then a suitable choice for $\Omega_P(J)$ is the singleton set consisting of the pair (J', ι) , where J' consists of the facts $R_1(a_1), R_2(a_2)$ and $\iota(a_1) = \iota(a_2) = a$. Indeed, for each \mathbf{S}_{in}^P -instance I , $P(I) \rightarrow J$ iff $I \rightarrow J'$ via homomorphisms that make the diagram in Definition 4.3 commute. Note that, to make the diagram commute, J' must indeed contain facts of the form $R_1(a_1)$ and $R_2(a_2)$ for distinct values a_1, a_2 that are both mapped to a by ι .

Our main result in this section is:

► **Theorem 4.5.** Every TAM Datalog program P has a generalized right-adjoint Ω_P . If P is connected, then $\Omega_P(J)$ is a singleton for all $J \in \text{Inst}[\mathbf{S}_{out}^P]$. Moreover, $\Omega_P(J)$ is computable from J and P in 2Exptime , and in ExpTime whenever the arity of P is bounded.

Proof. We first consider the special case of connected programs. We may assume without loss of generality that P is simple. This means that every rule is of the following form:

$$R_0(\mathbf{x}_0) :- E(\mathbf{y}), R_1(\mathbf{x}_1), \dots, R_m(\mathbf{x}_m) \tag{Eq. 2}$$

$$R(\mathbf{x}) :- E(\mathbf{y}), R_1(\mathbf{x}_1), \dots, R_m(\mathbf{x}_m) \tag{Eq. 3}$$

where E is an input relation, each R_i is an auxiliary relation, and R is an output relation.

To simplify the exposition below, we introduce some further notation. For each atom $R_i(\mathbf{x}_i)$ as in the above rule types, we will denote by $p_i \in \{1, \dots, n\}$ (with $n = \text{arity}(E)$) the unique number such that y_{p_i} is equal to the articulated variable in $R_i(\mathbf{x}_i)$. It indeed follows from the definition of TAM Datalog and the assumed connectedness and simplicity of P that such an index exists and is unique.

² By this we, we mean that, for all $a \in \text{adom}(I)$, either $h(a) = \iota(h'(a))$ or both are undefined.

10:10 Right-Adjoint for Datalog Programs

We construct an \mathbf{S}_{in} -instance J' consisting of all facts $E((b_1, X_1), \dots, (b_n, X_n))$ where

1. Each b_i is an element of $\text{adom}(J) \cup \{\perp\}$ and X_i is a set of \mathbf{S}_{aux} -facts over $\text{adom}(J) \cup \{\perp\}$ (not necessarily facts of J) in which b_i occurs in articulation position;
2. For each rule of the form (1) above and for each map $g : \{\mathbf{y}, \mathbf{x}_1, \dots, \mathbf{x}_m\} \rightarrow \text{adom}(J) \cup \{\perp\}$, if for each $1 \leq i \leq m$, $R_i(g(\mathbf{x}_i)) \in X_{p_i}$ then $R_0(g(\mathbf{x}_0)) \in X_{p_0}$; and
3. For each rule of the form (2) above and for each map $g : \{\mathbf{y}, \mathbf{x}_1, \dots, \mathbf{x}_m\} \rightarrow \text{adom}(J) \cup \{\perp\}$, if for each $1 \leq i \leq m$, $R_i(g(\mathbf{x}_i)) \in X_{p_i}$ then $R(g(\mathbf{x}))$ is a fact of J .

Note that the total number of possible facts $E((b_1, X_1), \dots, (b_n, X_n))$ in J' is double exponential in the combined size of P and J , and is exponential in J if the arity of P is bounded.

Let ι be the natural projection from J' to J , mapping all elements of the form (a, X) to a (and undefined on elements of the form (\perp, X)).

▷ **Claim.** For all \mathbf{S}_{in} -instances I , $P(I) \rightarrow J$ iff $I \rightarrow J'$. Moreover, the witnessing homomorphisms can be constructed so that the following diagram commutes:

$$\begin{array}{ccc} P(I) & \longrightarrow & J \\ id \uparrow & & \uparrow \iota \\ I & \longrightarrow & J' \end{array}$$

Proof. $[\Rightarrow]$ Let $h : P(I) \rightarrow J$. Recall that we denote by $\text{chase}_P(I)$ the $\mathbf{S}_{in} \cup \mathbf{S}_{out} \cup \mathbf{S}_{aux}$ -instance that is the chase of I (and of which I and $P(I)$ are the \mathbf{S}_{in} -reduct and \mathbf{S}_{out} -reduct, respectively). We extend h to the entire active domain of $\text{chase}_P(I)$ by sending every element a that is not in $\text{adom}(P(I))$ to a fresh value \perp . With a slight abuse of notation, in what follows, we denote by h the extended map from $\text{adom}(\text{chase}_P(I))$ to $\text{adom}(J) \cup \{\perp\}$. For each $a \in \text{adom}(\text{chase}_P(I))$, let F_a be the set of all \mathbf{S}_{aux} -facts of $\text{chase}_P(I)$ in which a occurs in articulation position. We define $h'(a) = (h(a), h(F_a))$.

We claim that h' is a homomorphism from I to J' . Let $E(a_1, \dots, a_n)$ be any fact of I . We must show that the fact $E((h(a_1), h(F_{a_1})), \dots, (h(a_n), h(F_{a_n})))$ belongs to J' . That is, we must show that conditions 1–3 hold.

Clearly, the first requirement is satisfied, namely, $h(X_{a_i})$ consists of facts in which $h(a_i)$ occurs in articulation position.

To see that the second requirement holds, consider a rule of form (1) and any map $g : \{\mathbf{y}, \mathbf{x}_1, \dots, \mathbf{x}_m\} \rightarrow \text{adom}(J)$, such that, for each $1 \leq i \leq m$, $R_i(g(\mathbf{x}_i)) \in X_{p_i}$. By construction, this means that each fact $R_i(g(\mathbf{x}_i))$ is the h -image of a fact $R_i(\mathbf{b}_i)$ in $\text{chase}_P(I)$. Now consider the map $\tilde{g} : \{\mathbf{y}, \mathbf{x}_1, \dots, \mathbf{x}_m\} \rightarrow \text{adom}(I)$ defined by $\tilde{g}(\mathbf{y}) = (a_1, \dots, a_n)$ and $\tilde{g}(\mathbf{x}_i) = \mathbf{b}_i$ for $i \leq m$. Note that \tilde{g} is a well-defined function. Indeed, for every $i \leq m$ and every $z \in R(\mathbf{x}_i)$, if z occurs more than once in the rule, then z must necessarily be the variable that appears in the articulation position p_i of $R(\mathbf{x}_i)$. It follows that z occurs in \mathbf{y} at position p_i , and, hence, necessarily, $R_i(g(\mathbf{x}_i))$ belongs to the h -image of $X_{a_{p_i}}$, and, hence, $\tilde{g}(z) = a_{p_i}$.

Since $\text{chase}_P(I)$ is closed under the rules of P , we may conclude that $R_0(\tilde{g}(\mathbf{x}_0))$ belongs to $\text{chase}_P(I)$. Let z be the variable occurring in articulation position in $R_0(\mathbf{x}_0)$. Recall that z occurs in \mathbf{y} at position p_0 , and $\tilde{g}(z) = a_{p_0}$. Then $R_0(\tilde{g}(\mathbf{x}_0)) \in F_{a_{p_0}}$, and hence, $h(F_{a_{p_0}})$ contains $R_0(h \circ \tilde{g}(\mathbf{x}_0))$. Note that by definition, $h \circ \tilde{g} = g$. In particular, $R_0(h \circ \tilde{g}(\mathbf{x}_0)) = R_0(g(\mathbf{x}_0))$. Therefore we have that $R_0(g(\mathbf{x}_0)) \in h(F_{a_{p_0}})$, and we are done.

To see that the third requirement holds, consider a rule of form (2) and any map $g : \{\mathbf{y}, \mathbf{x}_1, \dots, \mathbf{x}_m\} \rightarrow \text{adom}(J)$, such that, for each $1 \leq i \leq m$, $R_i(g(\mathbf{x}_i)) \in X_{p_i}$. By exactly the same reasoning as before, an h -preimage of the rule head $R(g(\mathbf{x}))$ belongs to $\text{chase}_P(I)$. Hence, it belongs to $P(I)$, therefore, $R(g(\mathbf{x}))$ is a fact of J .

It is also clear from the construction that $h \circ id = \iota \circ h'$, where id is the identity function on $\text{adom}(I) \cap \text{adom}(P(I))$. That is, the diagram commutes.

[\Leftarrow] Conversely, let $h : I \rightarrow J'$. Note that $\text{adom}(\text{chase}_P(I)) = \text{adom}(I)$, and hence we $h(a)$ is well-defined for all $a \in \text{adom}(\text{chase}_P(I))$. Let $h' : \text{adom}(I) \rightarrow \text{adom}(J)$ be the map such that $h'(a) = b$ whenever $h(a) = (b, X)$.

▷ **Subclaim 1.** For all $a \in \text{adom}(\text{chase}_P(I))$, if $h(a) = (b, X)$, then the h' -image of every \mathbf{S}_{aux} -fact of $\text{chase}_P(I)$ in which a occurs in articulation position belongs to X .

▷ **Subclaim 2.** h' is a homomorphism from $P(I)$ to J .

Subclaim 1 can be proved by induction on the derivation length of the fact in question.

To prove subclaim 2, let $R(\mathbf{a})$ be an \mathbf{S}_{out} -fact belonging to $P(I)$. Its derivation must use a rule of the form (2) above, using an assignment g (where $g(\mathbf{x}) = \mathbf{a}$). By Subclaim 1, we have that $R_i(h'(g(\mathbf{x}_i)))$ belongs to $h(g(y_{p_i}))_2$, for y_{p_i} the articulated variable in \mathbf{x}_i . Furthermore, $E(g(\mathbf{y}))$ holds in I , and hence $E(h(g(\mathbf{y})))$ holds in J' . By construction of J' , this means that the $R(h'(g(\mathbf{x})))$, that is, $R(h'(\mathbf{a}))$, belongs to J . This concludes the proof for the case of *connected* TAM Datalog programs.

It is also clear from the construction that $\iota \circ h = h' \circ id$, where id is the identity function on $\text{adom}(I) \cap \text{adom}(P(I))$. That is, the diagram commutes. \triangleleft

Finally, we show how to handle non-connected TAM Datalog programs. Let P be a non-connected TAM Datalog program. Let P' be obtained from P by adding a fresh binary input-relation S , and using this relation to make every every rule connected in some arbitrary way (more precisely, whenever the incidence graph of a rule body has multiple connected component, we add S -atoms to the body connecting these components while preserving tree-shapedness and almost-monadicity. For every input instance I , we denote by \hat{I} the $\mathbf{S}_{in} \cup \{S\}$ -instance extending I with all facts of the form $S(a, b)$ for $a, b \in \text{adom}(I)$. Furthermore, given an instance J' over the schema $\mathbf{S}_{in} \cup \{S\}$, by an “ S -component” of J' we will mean the \mathbf{S}_{in} -retract of a fully S -connected sub-instance of J' . Clearly, if J is an \mathbf{S}_{in} -instance and J' is a $\mathbf{S}_{in} \cup \{S\}$ -instance, then $\hat{J} \rightarrow J'$ iff $J \rightarrow J''$ for some S -component J'' of J' . Now we simply define $\Omega_P(J)$ to be the set of all S -components of instance in $\Omega_{P'}(J)$. Then we have: $P(I) \rightarrow J$ iff $P'(\hat{I}) \rightarrow J$ iff $\hat{I} \rightarrow \Omega_{P'}(J)$ iff $I \rightarrow J'$ for some $J' \in \Omega_P(J)$.

As a side remark, we mention that there is another way present the final argument where we lift the connected case to the general case: we can view the function that sends I to \hat{I} as a functor that itself has a generalized right-adjoint (sending I to its S -connected components). Thus, we can argue by composition of adjoints. \blacktriangleleft

We note that the proof of `thm:tam-adjoint` makes crucial use of both the tree-shapedness and the almost-monadicity of the Datalog program. Indeed, both properties are important for the existence of generalized right-adjoints as the following two propositions show:

► **Proposition 4.6.** *The tree-shaped Datalog program P in Example 3.11 (which is not almost-monadic) does not admit a generalized right-adjoint.*

Proof. Assume towards a contradiction that P has a generalized right-adjoint Ω_P . Let J be the two-element $\{R\}$ -instance consisting of the facts $R(0, 1)$ and $R(1, 0)$.

For $n \geq 1$, let C_n be the $\{E, F\}$ -instance consisting of the facts $E(v_0, v_1), \dots, E(v_{n-1}, v_n), E(v_n, v_0)$, that is, the directed E -cycle of length n . Trivially, $P(C_n) \rightarrow J$ for all $n \geq 1$. Therefore, for each $n \geq 1$, we have $C_n \rightarrow J'$ for some $J' \in \Omega_P(J)$. For every $J' \in \Omega_P(J)$ and for each element b of J' , let us define $n_{J',b}$ to be an arbitrarily chosen value such that

10:12 Right-Adjoints for Datalog Programs

$(C_n, v_0) \rightarrow (J', b)$, or undefined, if no such value exists. It follows from our earlier observation that $n_{J',b}$ is defined for at least one pair (J', b) with $J' \in \Omega_P(J)$. Let m be a common multiple of all defined $n_{J',b}$'s.

For each pair of positive integers $e \leq f$, let $I_{e,f}$ be the $\{E, F\}$ -instance depicted as follows:

$$u_0 \xrightarrow{E} v_0 \xrightarrow{F} u_1 \xrightarrow{E} v_1 \xrightarrow{F} u_2 \xrightarrow[\text{sequence of } e \text{ } E\text{-edges}]{\text{sequence of } f \text{ } F\text{-edges}} z \xrightarrow[\text{sequence of } f \text{ } F\text{-edges}]{\text{sequence of } e \text{ } E\text{-edges}} u_0$$

▷ **Claim 1.** For all $1 \leq e \leq f$, the following are equivalent:

1. $I_{e,f} \rightarrow J'$ for some $J' \in \Omega_P(J)$
2. $e \neq f$.

Proof. If $e = f$ then $P(I_{e,f})$ contains an R -cycle of odd length, viz. $u_0 \xrightarrow{R} u_1 \xrightarrow{R} u_2 \xrightarrow{R} u_0$, and therefore $P(I_{e,f}) \not\rightarrow J$. Hence, $I_{e,f} \not\rightarrow J'$ for all $J' \in \Omega_P(J)$. On the other hand, if $e < f$, then $P(I_{e,f})$ is a disjoint union of R -paths, and, clearly, $P(I_{e,f}) \rightarrow J$. Therefore, $I_{e,f} \rightarrow J'$ for some $J' \in \Omega_P(J)$. ◀

Now, let e be larger than the universe of all instances in $\Omega_P(J)$ and let $f = e + m$. By Claim 1, there is a homomorphism $h : I_{e,f} \rightarrow J'$ for some $J' \in \Omega_P(J)$. We will show that h can be extended to a homomorphism $h' : I_{f,f} \rightarrow J'$, which contradicts Claim 1. Let

$$u_2 = x_0 \xrightarrow{E} x_1 \cdots \xrightarrow{E} x_e = z$$

be the sub-instance of $I_{e,f}$ consisting of the E -edges joining u_2 and z . Similarly, let

$$u_2 = x_0 \xrightarrow{E} x_1 \cdots \xrightarrow{E} x_e \xrightarrow{E} x_{e+1} \cdots \xrightarrow{E} x_f = z$$

be the sub-instance of $I_{f,f}$ consisting of the E -edges joining u_2 and z . Recall that $f = e + m$. Since e is larger than the domain size of J' , it must be the case that $h(x_i) = h(x_j) = b$ for some $i < j \leq e$, and for some element b of J' . This means that b lies on a directed E -cycle in J' , and hence, in particular, it lies on a directed E -cycle of length m , say, $b = b_0 \xrightarrow{E} b_1 \cdots \xrightarrow{E} b_m = b$. The mapping $h' : I_{f,f} \rightarrow J'$ can be constructed simply by extending h and mapping x_{e+i} to b_i for $1 \leq i \leq m$. ◀

► **Proposition 4.7.** *The monadic Datalog program given by the single rule $\text{Ans}(x) :- E(x, x)$ does not admit a generalized right-adjoint.*

Proof. Let P be the Boolean Datalog program in question. Note that $\text{Unfoldings}(P, \text{Ans})$ consists of a pointed structure that is c-acyclic but not acyclic. It does not admit a generalized right-adjoint: let J be the empty \mathbf{S}_{out}^P -instance, and suppose for the sake of contradiction that there is a finite set $\{J_1, \dots, J_n\}$ such that, for all \mathbf{S}_{in}^P -instances I , $P(I) \rightarrow J$ iff $I \rightarrow J_i$ for some $i \leq n$. Let I_c be the instance consisting of a single reflexive Ans -edge of the form $\text{Ans}(a, a)$. Clearly, $P(I_c) \not\rightarrow J$, and therefore, $I_c \not\rightarrow J_i$. That is, J_1, \dots, J_n do not contain a reflexive Ans -edge. Next, let I_n be the instance that is an (irreflexive) Ans -clique of size n , where n is any number greater than the size of each J_i . Then, $I_n \not\rightarrow J_i$ (because if there was a homomorphism, J_i would contain a reflexive Ans -edge), but, trivially, $P(I_n) \rightarrow J$. ◀

In the special case of Boolean non-recursive programs, there is a converse to Thm. 4.5:

► **Theorem 4.8.** *A Boolean non-recursive Datalog program has a generalized right-adjoint iff it is equivalent to a TAM Datalog program.*

This follows from results that we will prove in Section 5 (specifically, Thm. 5.4 together with Thm. 5.3). It also follows from a known result about Pultr functors, namely [17, Theorem 2.5], since Boolean non-recursive Datalog programs define Pultr functors.

► **Remark 4.9.** Thm. 4.8 does not hold for *recursive* Boolean Datalog programs. Indeed, consider the Boolean Datalog program with input schema $\{R\}$ consisting of the rules

$$\begin{aligned} \text{Ans}() & \quad :- \quad \text{OddLengthPath}(x, x) \\ \text{OddLengthPath}(x, y) & \quad :- \quad R(x, y) \\ \text{OddLengthPath}(x, y) & \quad :- \quad R(x, z), R(z, u), \text{OddLengthPath}(u, y) \end{aligned}$$

It follows from well-known results in the CSP literature, (combined with Thm. 5.7 below), that P admits a generalized right-adjoint and that P is not equivalent to a monadic Datalog program. It follows by Lem. 3.8 that P is not equivalent to a TAM Datalog program.

► **Remark 4.10.** We note that the right adjoint $\Omega_P(\cdot)$ of a TAM Datalog program P is in general not definable by a Datalog program. This follows trivially from the fact that $\Omega_P(J)$ might consist of more than one structure if P is non connected and, in addition, It is not definable by a Datalog program even when P is connected as, in general, the domains of J and $\Omega_P(J)$ do not need to be related. For instance, Example 4.4 contains an example where the domain of $\Omega_P(J)$ is necessarily larger than that of J .

► **Remark 4.11.** Our definition of right-adjoints treats Datalog programs as functors in a flat category, while Lemma 2.3 naturally allows us to view Datalog programs as functors even in the non-flat category of instances and homomorphisms. This naturally raises a question of which functors between the “non-flat” categories allow right adjoints. This question has been answered by Pultr [22] up to some technical details, who described pairs of adjoint functors between categories of relational structures. Using either Pultr’s description, or by simple categorical methods, it can be seen that a Datalog program will rarely allow a proper right adjoint.

► **Remark 4.12.** While we are specifically interested in right-adjoints in this paper, one may also wonder what it means for a Datalog program to admit a (generalized) left-adjoint. Generalized left-adjoints for Datalog programs are closely related to query rewritings, as studied in the literature on data integration and data exchange. A Datalog program P has a generalized left-adjoint iff P is equivalent to a non-recursive Datalog program. Indeed, if P has a generalized left-adjoint Θ , then, for each $R \in \mathbf{S}_{out}^P$, the \mathbf{S}_{in}^P -instances in $\Theta(\{R(a_1, \dots, a_n)\})$ correspond to the members of $\text{Unfoldings}(P, R)$ (cf. [17, 13]).

5 Right Adjoints and Homomorphism Dualities

For any set of instances X , let $X^\uparrow = \{A \mid B \rightarrow A \text{ for some } B \in X\}$, and let $X^\downarrow = \{A \mid A \rightarrow B \text{ for some } B \in X\}$. A *homomorphism duality* is a pair of sets of instances (F, D) , such that F^\uparrow is the complement of D^\downarrow . The same definition extends naturally to pointed instances. By a *finite homomorphism duality*, we mean a homomorphism duality (F, D) where F and D are finite sets. By a *tree duality*, we mean a homomorphism duality (F, D) where F is a (possibly infinite) set of (not-necessarily-connected) acyclic instances, and D is finite.

The study of such dualities originated in combinatorics (see [20]) motivated by its links to the structure of the homomorphism partial order, and the complexity of deciding the existence of homomorphism between graphs and, more generally, relational structures (a.k.a. constraint satisfaction problems or CSPs). Indeed, dualities have played an important role in the study of CSPs. In particular, it was shown [3] that the CSPs definable in FO are precisely those

whose template is the right-hand side of a finite duality. In a similar vein, the CSPs solvable by the well-known *arc-consistency* algorithm are precisely those whose template is the right-hand side of a tree duality. More generally, the CSPs solvable by local consistency methods are those whose template is the right-hand side of a homomorphism duality whose left-hand side consists of instances of bounded treewidth. See [8] for a survey on the connections between duality and consistency algorithms. In database theory, homomorphism dualities are used in the study of the unique characterizability and exact learnability of schema mappings and database queries [2, 9], closed-world rewritings of open-world queries [6], and extremal fitting algorithms [10].

► **Example 5.1.** Let $\mathbf{S} = \{R\}$, where R is a binary relation symbol, and let $n \geq 1$. Let L_n be the finite linear order of length n , and let P_{n+1} be the directed path of length $n + 1$. Then $(\{P_{n+1}\}, \{L_n\})$ is a finite homomorphism duality.

► **Example 5.2.** Let $\mathbf{S} = \{P_0, P_1, E\}$, where P_0 and P_1 are unary and E is binary, and consider the two-element \mathbf{S} -instance $I = \{P_0(0), P_1(1), E(0, 0), E(1, 1)\}$ (without distinguished elements). For all \mathbf{S} -instances J , $J \rightarrow I$ holds if and only if no connected component of J contains both a P_0 -fact and a P_1 -fact. This can be expressed in the form of a tree duality: let F be the set of all (acyclic) instances consisting of an oriented path that connects a P_0 -node to a P_1 -node (where, by an *oriented path* we mean a path a_1, a_2, \dots, a_n where, for each $1 \leq i < n$, either $E(a_i, a_{i+1})$ or $E(a_{i+1}, a_i)$). Then $(F, \{I\})$ is a homomorphism duality.

► **Theorem 5.3** ([16, 9]). *Fix a schema \mathbf{S} and $k \geq 0$. Let F be any finite set of pairwise homomorphically incomparable k -ary pointed instances over \mathbf{S} . The following are equivalent:*

1. *There is a finite set of k -ary pointed instances D over \mathbf{S} such that (F, D) is a homomorphism duality.*
2. *Each pointed instance in F is homomorphically equivalent to a c -acyclic pointed instance. Moreover (for fixed \mathbf{S} and k), given a set F of c -acyclic pointed instances, such a set D can be computed in ExpTime .³*

The following theorem establishes a close relationship between generalized right-adjoints and homomorphism dualities:⁴

► **Theorem 5.4.** *Let P be any Datalog program that has a generalized right-adjoint. Then, for each $R \in \mathbf{S}_{out}^P$, there is a finite set of pointed \mathbf{S}_{in} -instances D such that $(\text{Unfoldings}(P, R), D)$ is a homomorphism duality.*

Proof. We may assume without loss of generality that $\mathbf{S}_{out}^P = \{R\}$. Let J be the \mathbf{S}_{out}^P -instance with $\text{adom}(J) = \{b_1, \dots, b_k, c\}$ (for $k = \text{arity}(R)$) containing all R -facts over $\text{adom}(J)$ except $R(b_1, \dots, b_k)$. Let $D = \{(J', \mathbf{b}') \mid (J', \iota) \in \Omega_P(J), \mathbf{b}' \in \text{adom}(J')^k, \iota(\mathbf{b}') = \mathbf{b}\}$, where $\mathbf{b} = b_1, \dots, b_k$. We claim that $(\text{Unfoldings}(P, R), D)$ is a homomorphism duality. Let (C, \mathbf{c}) be any \mathbf{S}_{in} -instance with k distinguished elements. Then an instance in $\text{Unfoldings}(P, R)$ homomorphically maps to (C, \mathbf{c}) iff $R(\mathbf{c}) \in P(C)$ iff $(P(C), \mathbf{c}) \not\rightarrow (J, \mathbf{b})$ iff (by the adjoint property) $(C, \mathbf{c}) \not\rightarrow (J', \mathbf{b}')$ for all $(J', \iota) \in \Omega_P(J)$ and \mathbf{b}' with $\iota(\mathbf{b}') = \mathbf{b}$. ◀

Thm. 5.4 shows that generalized right-adjoints can be used to construct duals. This approach was first used in [17] and [13], where right-adjoints of Pultr functors are applied to derive the dual of a tree. Thm. 5.4 can be viewed as extending results in [13] to a more general class of functors defined by recursive Datalog programs.

³ The ExpTime bound is not explicitly stated in [9] but follows from results in that paper.

⁴ Recall that $\text{Unfoldings}(P, R)$ can be viewed as an infinite union of (canonical instances of) conjunctive queries that defines the output relation R (cf. Lemma 2.5).

For TAM Datalog programs P , the proof of Thm. 5.4 yields a ExpTime algorithm to compute D from (P, R) provided the arity of the relations in P is bounded. Since $\text{Unfoldings}(P, R)$ consists of acyclic instances whenever P is a TAM Datalog program, this gives us a systematic way of constructing tree-dualities. In fact, every tree-duality can be obtained in this way:

► **Corollary 5.5.** *Let F be any set of acyclic pointed instances. The following are equivalent:*

1. *There is a finite set of pointed instances D such that (F, D) is a homomorphism duality*
2. *$F^\uparrow = \text{Unfoldings}(P, R)^\uparrow$ for some TAM Datalog program P and $R \in \mathbf{S}_{out}^P$.*

Proof. From 1 to 2: It is well known that, for any finite set of pointed instances D , there is an MSO formula ϕ that defines D^\downarrow . Hence, by duality, $\neg\phi$ defines F^\uparrow . Furthermore, the fact that F consists of acyclic pointed instances implies that $\neg\phi$ is tree-determined. Therefore, the direction 1 to 2 follows from Thm. 3.13. The direction from 2 to 1 follows from Thm. 5.4. ◀

It is possible to strengthen Corollary 5.5 by showing that the above conditions (1) and (2) are, in turn, equivalent to the fact that $F^\uparrow = G^\uparrow$ for some regular set G of acyclic queries (where “regular” needs to be defined in a suitable way, as in [15]). This follows from the fact that Thm. 3.13 uses tree-automata as an intermediate step in the proof. We note that the special case of this equivalence for Boolean CQs over digraphs was proven in [15].

Thm. 5.4 also implies that every finite set of acyclic pointed instances F is the left-hand side of a finite homomorphism duality: it suffices to let P be the TAM Datalog program containing a single non-recursive rule for each $(I, \mathbf{a}) \in F$, whose canonical instance is (I, \mathbf{a}) . Then, the unfoldings of P are, up to isomorphism, precisely the pointed instances in F . It follows from Thm. 5.4 that there is a finite set D such that (F, D) is a homomorphism duality. This provides an alternative proof of the characterization of left-hand sides of finite dualities given in [16] (i.e., the special case of Thm. 5.3 for structures without distinguished elements).

► **Remark 5.6.** In light of Thm. 5.3, it is natural to ask whether the above “dualities through adjoints” technique can be used to construct a finite homomorphism duality for any c-acyclic pointed instance. Prop. 4.7 shows that this is not possible. Note that the canonical instance of the rule of the program in Prop. 4.7 is $(\{E(a, a)\}, a)$, which is c-acyclic (but not acyclic).

For Boolean Datalog programs, the relationship between adjoints and dualities is tighter:

► **Theorem 5.7.** *For Boolean Datalog programs P , the following are equivalent:*

1. *P admits a generalized right-adjoint,*
2. *There is a finite set of pointed \mathbf{S}_{in}^P -instances D such that $(\text{Unfoldings}(P, \mathbf{Ans}), D)$ is a homomorphism duality.*

Proof. For every \mathbf{S}_{in}^P -instance I , $P(I)$ is either the empty instance, which we may denote as J_0 or the instance consisting of the zero-ary fact $\mathbf{Ans}()$, which we may denote as J_1 . Let $\Omega_P(J_0)$ be the set of all pairs (J', ι) with $J' \in D$ and ι the empty partial function; and let $\Omega_P(J_1) = \{(J', \iota)\}$ where J' is a single-element fully-connected \mathbf{S}_{in}^P -instance and ι is the empty partial function. It is easy to see that Ω_P is then a generalized right-adjoint for P . ◀

6 An Application: Generating Data Examples for Database Queries

We will now show-case one application of our results, which is concerned with the problem of generating data examples for database queries. A *data example* for a database query, informally, consists of a database instance I together with information about the output of the query when evaluated on I . Data examples can be a helpful tool in query debugging, query refinement, interactive query specification, and query learning. In each of these settings,

the question naturally comes up as to whether, for a given database query q , there exists a finite collection of data examples, such that, modulo logical equivalence, q is the *only* query (within some given class of queries) that fits the data examples. When this happens, we say that the collection of data examples in question *uniquely characterizes* the query q .

It was shown in [16, 9] that every “c-acyclic” union of conjunctive queries (UCQ) is indeed uniquely characterized by a finite collection of data examples. In fact, a UCQ is uniquely characterizable by a finite collection of data examples, if and only if it is equivalent to a c-acyclic UCQ. While this gives a precise answer to the question of unique characterizability, it can be cumbersome to use in practice. One of the reasons for this is that the data examples in question tend to look unnatural to a user. In particular, the existing algorithms for constructing data examples do not take integrity constraints into consideration. We will show here that the aforementioned results from [16, 9] can be adapted to the setting with integrity constraints, in such a way that all generated data examples satisfy the integrity constraints, provided that the integrity constraints are from a suitable, well-behaved class.

We will do this in three steps. First, we propose a suitable class of integrity constraints. Second, we study the existence of homomorphism dualities relative to a set of integrity constraints. Finally, we use this to construct uniquely characterizing examples for c-acyclic UCQs.

6.1 Tame sets of full TGDs

One of the most important classes of integrity constraints in databases is the class of *tuple-generating dependencies* (TGDs). The results we will present will be concerned with a subclass of TGDs called *full TGDs*. A full TGD is a TGD without existential quantifiers. More precisely, a full TGD is a first-order sentence of the form $\forall \mathbf{x}(\phi(\mathbf{x}) \rightarrow \psi(\mathbf{x}))$, where $\phi(\mathbf{x})$ and $\psi(\mathbf{x})$ are conjunctions of relational atomic formulas, and each variable in \mathbf{x} occurs in ϕ .

Every finite set of full TGDs naturally gives rise to a Datalog program. More precisely, for any set Σ of full TGDs over a schema \mathbf{S} , we will denote by P_Σ the Datalog program with $\mathbf{S}_{in}^P = \{R_{in} \mid R \in \mathbf{S}\}$, $\mathbf{S}_{out}^P = \{R_{out} \mid R \in \mathbf{S}\}$, and $\mathbf{S}_{aux}^P = \mathbf{S}$, consisting of the full TGDs in Σ as Datalog rules (where $\forall \mathbf{x}(\phi(\mathbf{x}) \rightarrow \psi(\mathbf{x}))$ becomes $\psi(\mathbf{x}) :- \phi(\mathbf{x})$), plus the “copy constraints” $R(\mathbf{x}) :- R_{in}(\mathbf{x})$ and $R_{out}(\mathbf{x}) :- R(\mathbf{x})$ for each $R \in \mathbf{S}$.

Although the input and output schemas of P_Σ are renamings of \mathbf{S} , we will write $P_\Sigma(I)$ even when I is an \mathbf{S} -instance instead of an \mathbf{S}_{in} , with the understanding that relation symbols are renamed in the obvious way; and similarly, we will freely treat the \mathbf{S}_{out} -instance $P_\Sigma(I)$ as an \mathbf{S} -instance. The Datalog program P_Σ then “captures” Σ in the following sense:

► **Lemma 6.1.** *Let Σ be any finite set of full TGDs. For all instances I , $P_\Sigma(I)$ is the unique minimal instance J with $I \subseteq J$ such that $J \models \Sigma$. In particular, $P_\Sigma(I) = I$ when $I \models \Sigma$.*

We say that a finite set Σ of full TGDs is *tame* (or, TAM-equivalent) if P_Σ is equivalent to a TAM Datalog program. A full TGD is *monadic* (*tree-shaped*) if the associated Datalog program P_Σ (where Σ consists of the full TGD in question) is monadic (resp. tree shaped).

► **Example 6.2.** The following sets of full TGDs are tame:

- $\Sigma_1 = \{\forall xyz(R(x, y) \wedge R(y, z) \rightarrow R(x, z))\}$. To see that P_{Σ_1} has a generalized right-adjoint, observe that it consists of the rules depicted on the left:

$$\begin{array}{ll}
 R(x, y) & :- R_{in}(x, y) \\
 R(x, z) & :- R(x, y), R(y, z) \\
 R_{out}(x, y) & :- R(x, y)
 \end{array}
 \qquad
 \begin{array}{ll}
 R(x, y) & :- R_{in}(x, y) \\
 R(x, z) & :- R(x, y), R_{in}(y, z) \\
 R_{out}(x, y) & :- R(x, y)
 \end{array}$$

P_{Σ_1} is not a TAM Datalog program. However, it is equivalent to the program P' consisting of the rules depicted on the right. Note how we have replaced one occurrence of R by R_{in} . The equivalence of P_{Σ_1} and P' is easy to show. Furthermore, P' is a TAM Datalog program (where the articulation position of R is the second position).

- $\Sigma_2 = \{\forall xyzu(R(x, y) \wedge R(y, z) \wedge R(z, u) \rightarrow R(x, u)), \forall xy(R(x, y) \rightarrow R(y, x))\}$. Although P_{Σ_2} is not a TAM Datalog program, it can be rewritten as one, using the same strategy as for Σ_1 . We will return to this example later, in Remark 6.6.
- Every finite set Σ of monadic tree-shaped TGDs. Indeed, P_Σ is a TAM Datalog program.

► **Remark 6.3.** The above example involves adhoc arguments. We leave it as an open problem to define a large syntactic class of (sets of) TGDs that have a generalized right-adjoint, which includes Σ_1 . Thm. 4.5 with Thm. 3.13 does imply that, for finite sets of tree-shaped TGDs Σ , if P_Σ is MSO-definable then Σ has a generalized right-adjoint.

Our main result in this section, namely Thm. 6.4, will apply to tame sets of full TGDs. The general strategy we develop here, however, can be extended to a larger class of integrity constraints that includes inclusion dependencies, as we will discuss in the conclusion section. This is beyond the scope of the present paper as it requires considering \exists Datalog programs (i.e., Datalog programs where existential quantifiers are allowed in rule heads).

6.2 Homomorphism dualities within restricted categories

Let \mathcal{C} be a class of instances over some schema (e.g., the class of transitive digraphs). We say that a pair (F, D) , with $F, D \subseteq \mathcal{C}$, is a *homomorphism duality within \mathcal{C}* if $F^\uparrow \cap \mathcal{C}$ is the complement of $D^\downarrow \cap \mathcal{C}$ relative to \mathcal{C} . In what follows we will also speak of homomorphism dualities *with respect to a theory Σ* . By this, we mean homomorphism dualities w.r.t. the class of instances defined by Σ . The next result shows how to obtain finite homomorphism dualities within \mathcal{C} , for classes \mathcal{C} that are definable by a tame set of full TGDs.

► **Theorem 6.4.** *Let Σ be a tame set of full TGDs. Let F be any finite set of pointed instances. If each member of F is of the form $(P_\Sigma(A), \mathbf{a})$ for some c-acyclic pointed instance (A, \mathbf{a}) , then F is the left-hand side of a finite duality w.r.t. Σ .*

Proof. Let F' be the finite set of c-acyclic pointed instances such that $F = \{(P_\Sigma(I), \mathbf{a}) \mid (I, \mathbf{a}) \in F'\}$. Let D be the finite set such that (F', D) is a homomorphism duality, given by Thm. 5.3. Let $D' = \{(P_\Sigma(B'), \mathbf{b}') \mid (B, \mathbf{b}) \in D, (B', \iota) \in \Omega_{P_\Sigma}(B), \iota(\mathbf{b}') = \mathbf{b}\}$. Note that D' consists of pointed instances satisfying Σ . We will show that (F, D') is a homomorphism duality w.r.t. Σ . Let (C, \mathbf{c}) be a pointed instance with $C \models \Sigma$. Then:

$$(C, \mathbf{c}) \in F^\uparrow \Leftrightarrow (C, \mathbf{c}) \in F'^\uparrow \Leftrightarrow (C, \mathbf{c}) \notin D^\downarrow \Leftrightarrow (C, \mathbf{c}) \notin D'^\downarrow$$

The first equivalence holds by Lem. 2.3 and the fact that $P_\Sigma(C) = C$. The second equivalence holds by the duality assumption. It remains to prove the third equivalence.

From left to right: By contraposition. Suppose that $(C, \mathbf{c}) \rightarrow (P_\Sigma(B'), \mathbf{b}')$ for some $(B, \mathbf{b}) \in D, (B', \iota) \in \Omega_{P_\Sigma}(B)$, and $\iota(\mathbf{b}') = \mathbf{b}$. Trivially, we have $id : (B', \mathbf{b}') \rightarrow (B', \mathbf{b}')$. It follows by the generalized adjoint property that $(P_\Sigma(B'), \mathbf{b}') \rightarrow (B, \iota(\mathbf{b}'))$. Therefore, by transitivity, and since $\iota(\mathbf{b}') = \mathbf{b}$, we have $(C, \mathbf{c}) \rightarrow (B, \mathbf{b})$ and therefore $(C, \mathbf{c}) \in D^\downarrow$.

From right to left: Again, by contraposition. Assume $(C, \mathbf{c}) \in D'^\downarrow$. Since $P_\Sigma(C) = C$, it follows that $(P_\Sigma(C), \mathbf{c}) \rightarrow (B, \mathbf{b})$ for some $(B, \mathbf{b}) \in D$. It follows by the adjoint property that $(C, \mathbf{c}) \rightarrow (B', \mathbf{b}')$ for some $(B, \mathbf{b}) \in D, (B', \iota) \in \Omega_{P_\Sigma}(B)$, and $\mathbf{b}' \in \iota^{-1}(\mathbf{b})$. Then also $(C, \mathbf{c}) \rightarrow (P_\Sigma(B'), \mathbf{b}')$. This means that $(C, \mathbf{c}) \in D'^\downarrow$. ◀

Regarding complexity, consider the case where Σ is a fixed tame set of full TGDs (not treated as part of the input). Then the proof of Thm. 6.4 yields a 2ExpTime algorithm for computing the dual set D from F , assuming F is specified by the underlying set of c -acyclic instances (A, \mathbf{a}) . Thus, for instance, for the class of transitive digraphs (which, as we saw earlier, is captured by a TAM Datalog program), we have a 2ExpTime-algorithm for constructing duals for digraphs that are specified as the transitive closure of an acyclic digraph.

The only prior results regarding homomorphism dualities for restricted classes of structures that we are aware of, are for undirected graphs and for finite algebras. An undirected graph can be viewed as an instance over a schema \mathbf{S} consisting of a single binary relation symbol E , satisfying the integrity constraints $\forall xy(E(x, y) \rightarrow E(y, x))$ and $\forall x \neg E(x, x)$. It is known that the category of undirected graphs and homomorphisms has no finite dualities, up to homomorphic equivalence, other than the trivial duality $(\{K_2\}, \{K_1\})$, where K_1 and K_2 are the 2-element clique and the empty graph, respectively (cf. [20]). Similarly, a finite algebra of a similarity type σ can be viewed as an \mathbf{S} -instance, with $\mathbf{S} = \{R_f \mid f \in \sigma\}$ satisfying $\Sigma = \{\forall \mathbf{x} \exists y R_f(\mathbf{x}, y), \forall \mathbf{x} y z (R_f(\mathbf{x}, y) \wedge R_f(\mathbf{x}, z) \rightarrow y = z) \mid f \in \sigma\}$, and, again, it is known that, in the category of finite algebras, no non-trivial finite dualities exist [4]. Note that both in the case of undirected graphs (viewed as symmetric and irreflexive relational structures) and in the case of finite algebras, all non-trivial structures in question are cyclic.

For the special case of monadic tree-shaped TGDs, we can prove a converse to Thm. 6.4:

► **Theorem 6.5.** *Let Σ be any set of monadic tree-shaped TGDs. Let F be any finite set of pairwise homomorphically-incomparable pointed instances (A, \mathbf{a}) with $A \models \Sigma$. Then, the following are equivalent:*

1. *F is the left hand side of a finite duality w.r.t. Σ ,*
2. *Each $(A, \mathbf{a}) \in F$ is homomorphically equivalent to $(P_\Sigma(A'), \mathbf{a})$ for some c -acyclic (A', \mathbf{a}) .*

► **Remark 6.6.** Thm. 6.5 cannot be lifted to arbitrary tame sets of full TGDs. Consider again the tame set of full TGDs $\Sigma = \{\forall xyz (R(x, y) \wedge R(y, z) \wedge R(z, u) \rightarrow R(x, u)), \forall xy (R(x, y) \rightarrow R(y, x))\}$ from Example 6.2. Let A be the instance (without distinguished elements) $\{R(a, a)\}$, and let B be the instance $\{R(a, b), R(b, a)\}$. Then $(\{A\}, \{B\})$ is a homomorphism duality w.r.t. Σ . Indeed, let C be an instance satisfying Σ and assume that $A \not\rightarrow C$ (i.e, C has no loop). Since C satisfies Σ it follows that C has no odd cycle and, hence, is homomorphic to B . However, it is easy to see that every instance A' satisfying $P_\Sigma(A') = A$ must have a cycle.

6.3 Uniquely Characterizing Examples for Database Queries

By a *collection of labeled examples* for a k -ary query, we mean a pair (E^+, E^-) of finite sets of pointed instances with k distinguished elements.⁵ A UCQ q fits such (E^+, E^-) if $\mathbf{a} \in q(A)$ for all $(A, \mathbf{a}) \in E^+$, and $\mathbf{a} \notin q(A)$ for all $(A, \mathbf{a}) \in E^-$. We say that two UCQs q, q' (over the same schema \mathbf{S}) are *equivalent w.r.t. Σ* , where Σ is a first-order theory, if for all $I \in \text{Inst}[\mathbf{S}]$ with $I \models \Sigma$, $q(I) = q'(I)$. A collection of labeled examples (E^+, E^-) *uniquely characterizes* a UCQ q w.r.t. Σ , if q fits (E^+, E^-) , and every UCQ that fits (E^+, E^-) is equivalent to q w.r.t. Σ .

⁵ Data examples can also be defined as pairs $(I, q(I))$ of an input instance and the complete query output. This would not change our results. Note that such a data example $(I, q(I))$ can be equivalently represented by all positive examples (I, \mathbf{a}) for $\mathbf{a} \in q(I)$ and negative examples (I, \mathbf{a}) for $\mathbf{a} \in \text{adom}(I)^k \setminus q(I)$.

► **Lemma 6.7.** *Let \mathbf{S} be any schema and Σ a FO theory over \mathbf{S} . Let q be a UCQ over \mathbf{S} , and let E^+, E^- be finite sets of pointed instances (I, \mathbf{a}) with $I \models \Sigma$. The following are equivalent:*

1. *The collection of labeled examples (E^+, E^-) uniquely characterizes q w.r.t. Σ*
2. *q fits (E^+, E^-) and (E^+, E^-) is a finite homomorphism duality w.r.t. Σ .*

► **Theorem 6.8.** *Let Σ a tame set of full TGDs over a schema \mathbf{S} . Every c-acyclic UCQ q over \mathbf{S} is uniquely characterized w.r.t. Σ by a collection of labeled examples satisfying Σ .*

Proof. Let E^+ be the set of pointed instances $(P_\Sigma(I), \mathbf{a})$, for (I, \mathbf{a}) a (c-acyclic) canonical instance of a CQ in q . By Thm. 6.4, there is a finite set E^- such that (E^+, E^-) is a homomorphism duality w.r.t. Σ . By Lem. 6.7, (E^+, E^-) uniquely characterizes q w.r.t. Σ . ◀

The proof of Thm. 6.8 yields a 2ExpTime upper bound for computing uniquely characterizing examples for a given c-acyclic UCQ, relative to a fixed tame set of full TGDs. It is not known whether this is optimal. In the absence of integrity constraints, uniquely characterizing examples for a c-acyclic UCQ can be constructed in ExpTime and this is known to be optimal since they are in general exponential in size. In the case of c-acyclic CQs, uniquely characterizing examples can be constructed in polynomial time [9].

7 Conclusion

We introduced a new fragment of Datalog, TAM Datalog, that is semantically well-behaved (closed under composition and having a natural semantic characterization) and admits generalized right-adjoints. We used this result to obtain a method for constructing uniquely characterizing data examples for c-acyclic UCQs in the presence of integrity constraints (where the data examples are required to satisfy the integrity constraints). Generalized right-adjoints for Datalog programs seem potentially useful in other contexts as well, such as in tasks involving reasoning about a hidden database instance based on an exposed view (cf. [5]).

In a companion paper (cf. [11]), we further extend our study to \exists Datalog (the extension of Datalog where existential quantifiers are allowed in rule heads), and we show that linear \exists Datalog programs have right-adjoints. This is then used to extend our results on uniquely characterizing data examples to the case with (a weakly acyclic set of) inclusion dependencies.

We leave as open problems for future research: (i) obtaining tight complexity bounds for the task of constructing uniquely characterizing data examples for CQs and UCQs in the presence of a tame set of full TGDs; (ii) identifying better syntactic criteria that guarantees that a given finite set of full TGDs is tame; (iii) extending our results to the case with functional dependencies⁶; and (iv) studying which logic (or, more generally, formalism) is necessary to define the right adjoint $\Sigma_P(\cdot)$ (see also Remark 4.10).

References

- 1 Serge Abiteboul, Richard Hull, and Victor Vianu. *Foundations of Databases*. Addison Wesley, 1995.
- 2 Bogdan Alexe, Balder ten Cate, Phokion G. Kolaitis, and Wang-Chiew Tan. Characterizing schema mappings via data examples. *ACM Trans. Database Syst.*, 36(4):23:1–23:48, dec 2011. doi:10.1145/2043652.2043656.

⁶ Recent results in [18] show that \mathcal{ELI} -concepts (equivalently: unary connected acyclic CQs over a schema with only binary relations) are uniquely characterizable in the presence of functional dependencies, and that uniquely characterizing examples can be computed for them in polynomial time.

- 3 Albert Atserias. On digraph coloring problems and treewidth duality. *Eur. J. Comb.*, 29(4):796–820, 2008. doi:10.1016/j.ejc.2007.11.004.
- 4 Richard N. Ball, Jaroslav Nešetřil, and Aleš Pultr. Dualities in full homomorphisms. *European Journal of Combinatorics*, 31(1):106–119, 2010. doi:10.1016/j.ejc.2009.04.004.
- 5 Michael Benedikt, Pierre Bourhis, Balder Ten Cate, Gabriele Puppis, and Michael Vanden Boom. Inference from visible information and background knowledge. *ACM Trans. Comput. Logic*, 22(2), jun 2021. doi:10.1145/3452919.
- 6 Meghyn Bienvenu, Balder Ten Cate, Carsten Lutz, and Frank Wolter. Ontology-based data access: A study through Disjunctive Datalog, CSP, and MMSNP. *ACM Trans. Database Syst.*, 39(4), dec 2015. doi:10.1145/2661643.
- 7 Manuel Bodirsky, Simon Knäuer, and Sebastian Rudolph. Datalog-expressibility for monadic and guarded second-order logic. In *ICALP 2021*, pages 120:1–120:17, 2021. doi:10.4230/LIPIcs.ICALP.2021.120.
- 8 Andrei A. Bulatov, Andrei A. Krokhin, and Benoît Larose. Dualities for constraint satisfaction problems. In Nadia Creignou, Phokion G. Kolaitis, and Heribert Vollmer, editors, *Complexity of Constraints - An Overview of Current Research Themes [Result of a Dagstuhl Seminar]*, volume 5250 of *LNCS*, pages 93–124. Springer, 2008. doi:10.1007/978-3-540-92800-3_5.
- 9 Balder ten Cate and Victor Dalmau. Conjunctive queries: Unique characterizations and exact learnability. *ACM Trans. Database Syst.*, 47(4), nov 2022. doi:10.1145/3559756.
- 10 Balder ten Cate, Victor Dalmau, Maurice Funk, and Carsten Lutz. Extremal fitting problems for conjunctive queries. In *Proceedings of the 42nd ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, PODS’23*, 2023. doi:10.1145/3584372.3588655.
- 11 Balder ten Cate, Víctor Dalmau, and Jakub Opršal. Right-adjoints for datalog programs, and homomorphism dualities over restricted classes, 2023. doi:10.48550/arXiv.2302.06366.
- 12 Surajit Chaudhuri and Moshe Y Vardi. On the equivalence of recursive and nonrecursive Datalog programs. *Journal of Computer and System Sciences*, 54(1):61–78, 1997. doi:10.1006/jcss.1997.1452.
- 13 Victor Dalmau, Andrei Krokhin, and Jakub Opršal. Functors on relational structures which admit both left and right adjoints, 2023. doi:10.48550/arXiv.2302.13657.
- 14 Víctor Dalmau and Jakub Opršal. Local consistency as a reduction between constraint satisfaction problems, 2023. doi:10.48550/arXiv.2301.05084.
- 15 Péter L. Erdős, Dömötör Pálvölgyi, Claude Tardif, and Gábor Tardos. Regular families of forests, antichains and duality pairs of relational structures. *Comb.*, 37(4):651–672, 2017. doi:10.1007/s00493-015-3003-4.
- 16 Jan Foniok, Jaroslav Nešetřil, and Claude Tardif. Generalised dualities and maximal finite antichains in the homomorphism order of relational structures. *Eur. J. Comb.*, 29(4):881–899, 2008. doi:10.1016/j.ejc.2007.11.017.
- 17 Jan Foniok and Claude Tardif. Digraph functors which admit both left and right adjoints. *Discrete Mathematics*, 338(4):527–535, 2015. doi:10.1016/j.disc.2014.10.018.
- 18 Maurice Funk, Jean Christoph Jung, and Carsten Lutz. Frontiers and exact learning of ELI queries under DL-Lite ontologies. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence, IJCAI-22*, 2022. doi:10.24963/ijcai.2022/364.
- 19 Georg Gottlob and Christoph Koch. Monadic Datalog and the expressive power of languages for web information extraction. *J. ACM*, 51(1):74–113, jan 2004. doi:10.1145/962446.962450.
- 20 Pavol Hell and Jaroslav Nešetřil. *Graphs and homomorphisms*, volume 28 of *Oxford lecture series in mathematics and its applications*. Oxford University Press, 2004.
- 21 Andrei A. Krokhin, Jakub Opršal, Marcin Wrochna, and Stanislav Živný. Topology and adjunction in promise constraint satisfaction. *SIAM Journal on Computing*, 52(1):38–79, 2023. doi:10.1137/20M1378223.
- 22 Aleš Pultr. The right adjoints into the category of relational systems. In *Reports of the Midwest Category Seminar IV*, volume 137 of *Lecture Notes in Mathematics*, pages 100–113. Springer, 1970.
- 23 Sebastian Rudolph and Markus Krötzsch. Flag & check: Data access with monadically defined queries. In *Proceedings of the 32nd ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, PODS ’13*, pages 151–162. ACM, 2013. doi:10.1145/2463664.2465227.