

BIROn - Birkbeck Institutional Research Online

Fortin, M. and Konev, B. and Ryzhikov, Vladislav and Savateev, Yury and Wolter, F. and Zakhariyashchev, Michael (2022) Unique characterisability and learnability of Temporal Instance Queries. In: 19th International Conference on Principles of Knowledge Representation and Reasoning, KR 2022, 31 Jul - 05 Aug 2022, Haifa, Israel.

Downloaded from: <https://eprints.bbk.ac.uk/id/eprint/48209/>

Usage Guidelines:

Please refer to usage guidelines at <https://eprints.bbk.ac.uk/policies.html>

or alternatively

contact lib-eprints@bbk.ac.uk.

Unique Characterisability and Learnability of Temporal Instance Queries

M. Fortin¹, B. Konev¹, V. Ryzhikov², Y. Savateev², F. Wolter¹, M. Zakharyashev²

¹Department of Computer Science, University of Liverpool, UK

²Department of Computer Science, Birkbeck, University of London, UK

{marie.fortin,boris.konev,wolter@liverpool.ac.uk}, {vlad,yury,michael}@dcs.bbk.ac.uk

Abstract

We aim to determine which temporal instance queries can be uniquely characterised by a (polynomial-size) set of positive and negative temporal data examples. We start by considering queries formulated in fragments of propositional linear temporal logic *LTL* that correspond to conjunctive queries (CQs) or extensions thereof induced by the until operator. Not all of these queries admit polynomial characterisations but by restricting them further to path-shaped queries we identify natural classes that do. We then investigate how far the obtained characterisations can be lifted to temporal knowledge graphs queried by 2D languages combining *LTL* with concepts in description logics \mathcal{EL} or \mathcal{ELI} (i.e., tree-shaped CQs). While temporal operators in the scope of description logic constructors can destroy polynomial characterisability, we obtain general transfer results for the case when description logic constructors are within the scope of temporal operators. Finally, we apply our characterisations to establish (polynomial) learnability of temporal instance queries using membership queries in the active learning framework.

1 Introduction

Constructing queries or, more generally, logical concepts describing individuals of interest, can be difficult. Providing support to a user to cope with this problem has been a major research topic in databases, logic, and knowledge representation. For instance, in reverse engineering of database queries and concept descriptions (Martins 2019; Lehmann and Hitzler 2010; Jung et al. 2020), one aims to identify a query using a set of positively and negatively labelled examples of answers and non-answers, respectively; and in active learning approaches, one aims to identify a query by asking an oracle (e.g., domain specialist) whether an example is an answer or a non-answer to the query (Angluin, Frazier, and Pitt 1992; Funk, Jung, and Lutz 2021; ten Cate and Dalmau 2021).

Recently, the *unique characterisation* of a query by a finite (ideally, polynomial-size) set of positive and negative example answers has been identified as a fundamental link between queries and data (ten Cate and Dalmau 2021). Namely, we say that a query q fits a pair $E = (E^+, E^-)$ of sets E^+ and E^- of pointed databases (\mathcal{D}, a) if $\mathcal{D} \models q(a)$ for $(\mathcal{D}, a) \in E^+$, and $\mathcal{D} \not\models q(a)$ for $(\mathcal{D}, a) \in E^-$. Then E uniquely characterises q within a class \mathcal{Q} of queries if q is the only (up to equivalence) query in \mathcal{Q} that fits E .

Unique (polynomial) characterisations can be used to illustrate, explain, and construct queries. They are also a ‘non-procedural’ necessary condition for (polynomial) learnability using membership queries in Angluin’s (1987b) framework of active learning, where membership queries to the oracle take the form ‘does $\mathcal{D} \models q(a)$ hold?’. It is shown by ten Cate and Dalmau (2021) that, for classes of conjunctive queries (CQs), it is often a sufficient condition as well.

In many applications, queries are required to capture the temporal evolution of individuals, making their formulation even harder. The aim of this paper is to start an investigation of the (polynomial) characterisability of temporal instance queries. We first consider one-dimensional data instances of the form $(\delta_0, \dots, \delta_n)$, where δ_i is the set of atomic propositions that are true at timestamp i , describing the temporal behaviour of a single individual, and queries formulated in fragments of propositional linear temporal logic *LTL*. Although rather basic as a temporal data model, this restriction allows us to focus on the purely temporal aspect of unique characterisability. We then generalise our results, where possible, to standard two-dimensional temporal data instances, in which the δ_i are replaced by non-temporal data instances and queries are obtained by combining fragments of *LTL* with \mathcal{ELI} -concept queries (or tree-shaped CQs), thereby combining a well established formalism for accessing temporal data (Chomicki and Toman 2018) with the basic concept descriptions for tractable data access from description logic (Baader et al. 2017).

Our initial observation is that already very primitive temporal queries are not uniquely characterisable. For example¹, consider the query $q = \diamond_r(A \wedge B)$ with the operator \diamond_r ‘now or later’ (interpreted by \leq over linearly ordered timestamps). By the pigeonhole principle, no finite example set E can distinguish q from a query $q' = \diamond_r(A \wedge (\diamond_r B \wedge \diamond_r(A \wedge \dots)))$ with sufficiently many alternating A and B . Similarly, the query $q = \circ A$ with the ‘next time’ operator \circ is not distinguishable by a finite example set from $q' = (\circ \dots \circ A) \cup A$ with the strict ‘until’ operator \cup and sufficiently many \circ on its left-hand side.

Aiming to identify natural and useful classes of temporal queries enjoying (polynomial) characterisability, in this

¹For detailed explanations and omitted proofs, the reader is referred to the full Arxiv version (Fortin et al. 2022b) of this paper.

paper we consider the conjunctive fragment of *LTL*. To begin with, we focus on two classes of path CQs: the class $\mathcal{Q}_p[\circ, \diamond_r]$ of queries of the form

$$\mathbf{q} = \rho_0 \wedge \mathbf{o}_1(\rho_1 \wedge \mathbf{o}_2(\rho_2 \wedge \dots \wedge \mathbf{o}_n \rho_n)), \quad (1)$$

where $\mathbf{o}_i \in \{\circ, \diamond_r\}$ and ρ_i is a conjunction of atomic propositions, and also the class $\mathcal{Q}_p^\sigma[\mathbf{U}]$ of U-queries of the form

$$\mathbf{q} = \rho_0 \wedge (\lambda_1 \mathbf{U} (\rho_1 \wedge (\lambda_2 \mathbf{U} (\dots (\lambda_n \mathbf{U} \rho_n) \dots))), \quad (2)$$

where λ_i is a conjunction of atoms or \perp . The superscript σ in $\mathcal{Q}_p^\sigma[\mathbf{U}]$ indicates that queries are formulated in a finite signature σ , a condition required because of the universal quantification implicit in U. Our first main result is a syntactic criterion of (polynomial) characterisability of $\mathcal{Q}_p[\circ, \diamond_r]$ -queries. In fact, it turns out that the query $\diamond_r(A \wedge B)$ mentioned above epitomises the cause of non-characterisability in $\mathcal{Q}_p[\circ, \diamond_r]$. It follows, in particular, that the restriction $\mathcal{Q}_p[\circ, \diamond]$ of $\mathcal{Q}_p[\circ, \diamond_r]$ to queries with \circ and strict eventuality $\diamond = \circ \diamond_r$ is polynomially characterisable. Our second main result is that all $\mathcal{Q}_p^\sigma[\mathbf{U}]$ -queries with \subseteq -incomparable λ_i and ρ_i , for each i , are polynomially characterisable within $\mathcal{Q}_p^\sigma[\mathbf{U}]$. Although we show that all $\mathcal{Q}_p^\sigma[\mathbf{U}]$ -queries are exponentially characterisable, it remains open whether they are polynomially characterisable in $\mathcal{Q}_p^\sigma[\mathbf{U}]$.

The essential property that distinguishes $\mathcal{Q}_p[\circ, \diamond_r]$ and $\mathcal{Q}_p^\sigma[\mathbf{U}]$ from other queries is that they do not admit temporal branching as, for instance, in $\diamond A \wedge \diamond B$. In fact, we show that even within the class of queries using only \wedge and \diamond and with a bound on the number of branches, not all queries are polynomially characterisable. A first step towards positive results covering non-path queries is made for the case of queries in which all branches are of equal length.

Our next aim is to generalise the obtained results to 2D temporal queries combining *LTL* with the description logic constructor $\exists P$ of \mathcal{ELI} . Our first main result is negative: even queries of the form $\exists P.q_1 \wedge \dots \wedge \exists P.q_n$, in which $q_i \in \mathcal{Q}_p[\circ, \diamond]$, are not polynomially characterisable. The situation changes drastically, however, if we consider queries of the form (1) or (2), in which ρ_i and λ_i are \mathcal{ELI} -queries. Indeed, we generalise our polynomial characterisability results for $\mathcal{Q}_p[\circ, \diamond_r]$ and $\mathcal{Q}_p^\sigma[\mathbf{U}]$ to such queries using recent results on the computation of frontiers in the lattice of \mathcal{ELI} -queries (ten Cate and Dalmau 2021) and proving a new result on split partners in the lattice of \mathcal{EL} -queries (where \mathcal{EL} denotes \mathcal{ELI} without inverse roles).

Finally, we discuss applications of our results to learning temporal instance queries using membership queries of the form ‘does $\mathcal{D} \models \mathbf{q}$ hold?’. As we always construct example sets effectively, our unique (exponential) characterisability results imply (exponential-time) learnability with membership queries. Obtaining polynomial-time learnability from polynomial characterisations is more challenging. A main result here is that $\mathcal{Q}_p[\circ, \diamond_r]$ with \mathcal{ELI} -queries is polynomial-time learnable with membership queries, assuming the learner is given the target query size in advance.

2 Related Work

Our contribution is closely related to work on active learnability of formal languages and on learning temporal logic

formulas interpreted over finite and infinite traces. It is also related to learning database queries and other formal expressions for accessing data. In the former area, the seminal paper by Angluin (1987a) has given rise to a large body of work on active learning of regular languages or variations, for example, (Shahbaz and Groz 2009; Aarts and Vaandrager 2010; Cassel et al. 2016; Howar and Steffen 2018). This work has mainly focused on learning various types of finite state machines or automata using a combination of membership queries with other powerful types of queries such as equivalence queries. The use of two or more types of queries is motivated by the fact that otherwise one cannot efficiently learn a wide variety of important languages, including regular languages. In fact, the main difference between this work and our contribution is that we focus on queries for which the corresponding formal languages form only a small subset of the regular languages and it is this restriction that enables us to focus on characterisability and learnability with membership queries.

Rather surprisingly, there has hardly been any work on active learning of temporal formulas over finite or infinite traces; we refer the reader to (Camacho and McIlraith 2019), also for a discussion of the relationship between learning automata and *LTL*-formulas. In contrast, passive learning of *LTL*-formulas has recently received significant attention; see (Lemieux, Park, and Beschastnikh 2015; Neider and Gavran 2018; Fijalkow and Lagarde 2021; Fortin et al. 2022a) and, in the context of explainable AI, also (Camacho and McIlraith 2019) for an overview.

In the database and KR communities, there has been extensive work on identifying queries and concept descriptions from data examples. For instance, in reverse engineering of queries, the goal is typically to decide whether there is a query that fits (or separates) a set of positive and negative examples. Relevant work under the closed world assumption include (Arenas and Diaz 2016; Barceló and Romero 2017) and under the open world assumption (Gutiérrez-Basulto, Jung, and Sabellek 2018; Funk et al. 2019). Related work on active learning not yet discussed include the identification of \mathcal{EL} -queries (Funk, Jung, and Lutz 2021) and ontologies (Konev, Ozaki, and Wolter 2016; Konev et al. 2017), and of schema-mappings (ten Cate, Dalmau, and Kolaitis 2013; ten Cate et al. 2018). Again this work differs from our contribution as it focuses on learning using membership and equivalence queries rather than only the former. The use of unique characterisations to explain and construct schema mappings has been promoted and investigated by Kolaitis (2011) and Alexe et al. (2011).

Combining *LTL* and description logics for temporal conceptual modelling and data access has a long tradition (Lutz, Wolter, and Zakharyashev 2008; Artale et al. 2017). For querying purposes, sometimes description logic concepts have been replaced by general CQs. Our restriction to \mathcal{ELI} -concepts instead of general CQs is motivated by results of (ten Cate and Dalmau 2021) showing that only CQs that are acyclic modulo cycles through the answer variables are polynomially characterisable within the class of CQs. Hence very strong acyclicity conditions are needed to ensure polynomial characterisability. We conjecture that our results can

be extended to this class.

The class of queries in which no $\exists P$ is within the scope of temporal operators was first introduced by (Baader, Borgwardt, and Lippmann 2015; Borgwardt and Thost 2015) in the context of monitoring applications. The lcs and msc in temporal DLs are considered by Tirtarasa and Turhan 2022.

3 Preliminaries

By a *signature* we mean any finite set $\sigma \neq \emptyset$ of *atomic concepts* A, B, C, \dots representing observations, measurements, events, etc. A σ -*data instance* is any finite sequence $\mathcal{D} = (\delta_0, \dots, \delta_n)$ with $\delta_i \subseteq \sigma$, saying that $A \in \delta_i$ happened at moment i . The *length* of \mathcal{D} is $\max(\mathcal{D}) = n$ and the *size* of \mathcal{D} is $|\mathcal{D}| = \sum_{i \leq n} |\delta_i|$. We do not distinguish between \mathcal{D} and its variants of the form $(\delta_0, \dots, \delta_n, \emptyset, \dots, \emptyset)$.

We access data by means of *queries*, q , constructed from atoms, \perp and \top using \wedge and the temporal operators $\circ, \diamond, \diamond_r$ and \cup . The set of atomic concepts occurring in q is denoted by $\text{sig}(q)$. The set of queries that only use the operators from $\Phi \subseteq \{\circ, \diamond, \diamond_r, \cup\}$ is denoted by $\mathcal{Q}[\Phi]$; $\mathcal{Q}^\sigma[\Phi]$ is the restriction of $\mathcal{Q}[\Phi]$ to a signature σ . The *size* $|q|$ of q is the number of symbols in q , and the *temporal depth* $\text{tdp}(q)$ of q is the maximum number of nested temporal operators in q .

$\mathcal{Q}[\circ, \diamond, \diamond_r]$ -queries can be equivalently defined as *tree-shaped* conjunctive queries (CQs) with the binary predicates $\text{succ}, <, \leq$ over \mathbb{N} , and atomic concepts as unary predicates. Each such CQ is a set $Q(t_0)$ of assertions of the form $A(t)$, $\text{succ}(t, t')$, $t < t'$, and $t \leq t'$, with a distinguished variable t_0 , such that, for every variable t in $Q(t_0)$, there exists exactly one path from t_0 to t along the binary predicates $\text{succ}, <, \leq$.

The set of $\mathcal{Q}[\circ, \diamond, \diamond_r]$ -queries with *path-shaped* CQ counterparts is denoted by $\mathcal{Q}_p[\circ, \diamond, \diamond_r]$. Such queries q take the form (1), where $\mathbf{o}_i \in \{\circ, \diamond, \diamond_r\}$ and ρ_i is a conjunction of atoms (the empty conjunction is \top). Similarly, $\mathcal{Q}_p[\cup]$ -queries take the form (2).

Given a data instance $\mathcal{D} = (\delta_0, \dots, \delta_n)$, the *truth-relation* $\mathcal{D}, \ell \models q$, for $\ell < \omega$, is defined as follows:

$$\begin{aligned} \mathcal{D}, \ell \models A &\text{ iff } A \in \delta_\ell, & \mathcal{D}, \ell \models \top, & \mathcal{D}, \ell \not\models \perp, \\ \mathcal{D}, \ell \models q_1 \wedge q_2 &\text{ iff } \mathcal{D}, \ell \models q_1 \text{ and } \mathcal{D}, \ell \models q_2, \\ \mathcal{D}, \ell \models \circ q &\text{ iff } \mathcal{D}, \ell + 1 \models q, \\ \mathcal{D}, \ell \models \diamond q &\text{ iff } \mathcal{D}, m \models q, \text{ for some } m > \ell, \\ \mathcal{D}, \ell \models \diamond_r q &\text{ iff } \mathcal{D}, m \models q, \text{ for some } m \geq \ell, \\ \mathcal{D}, \ell \models q_1 \cup q_2 &\text{ iff there is } m > \ell \text{ such that } \mathcal{D}, m \models q_2 \\ &\text{ and } \mathcal{D}, k \models q_1, \text{ for all } k \text{ with } \ell < k < m. \end{aligned}$$

Note that $\mathcal{D}, n \models \diamond \top \wedge \circ \top \wedge (q \cup \top)$ as $(\delta_0, \dots, \delta_n, \emptyset)$ is a variant of \mathcal{D} . We write $q \models q'$ if $\mathcal{D}, \ell \models q$ implies $\mathcal{D}, \ell \models q'$ for any \mathcal{D} and ℓ . If $q \models q'$ and $q' \models q$, we call q and q' *equivalent* and write $q \equiv q'$. Since $\circ q \equiv \perp \cup q$, $\diamond q \equiv \top \cup q$ and $\diamond q \equiv \circ \diamond_r q$, one can assume that $\mathcal{Q}[\circ, \diamond] \subseteq \mathcal{Q}[\cup]$, $\mathcal{Q}[\diamond] \subseteq \mathcal{Q}[\circ, \diamond_r]$ and $\mathcal{Q}[\circ, \diamond_r] = \mathcal{Q}[\circ, \diamond_r, \diamond]$.

4 Unique Characterisability

An *example set* is a pair $E = (E^+, E^-)$ with finite sets E^+ and E^- of data instances that are called *positive* and *negative examples*, respectively. A query q *fits* E if $\mathcal{D}^+, 0 \models q$

and $\mathcal{D}^-, 0 \not\models q$, for all $\mathcal{D}^+ \in E^+$ and $\mathcal{D}^- \in E^-$. We say that E *uniquely characterises* q within a class \mathcal{Q} of queries if q fits E and $q \equiv q'$ for any $q' \in \mathcal{Q}$ that fits E . If all $q \in \mathcal{Q}$ are characterised by some E within $\mathcal{Q}' \supseteq \mathcal{Q}$, we call \mathcal{Q} *uniquely characterisable* within \mathcal{Q}' . Further, \mathcal{Q} is *polynomially characterisable* within $\mathcal{Q}' \supseteq \mathcal{Q}$ if there is a polynomial f such that every $q \in \mathcal{Q}$ is characterised within \mathcal{Q}' by some E of size $|E| \leq f(|q|)$, where $|E| = \sum_{\mathcal{D} \in (E^+ \cup E^-)} |\mathcal{D}|$. Let \mathcal{Q}^n be the set of queries in \mathcal{Q} of size at most n . We say that \mathcal{Q} is *polynomially characterisable for bounded query size* if there is a polynomial f such that every $q \in \mathcal{Q}^n$ is characterised by some E of size $\leq f(n)$ within \mathcal{Q}^n .

Observe that (polynomial) characterisability is anti-monotone: if a query q is (polynomially) characterisable within \mathcal{Q} and $\mathcal{Q}' \subseteq \mathcal{Q}$, then q is (polynomially) characterisable within \mathcal{Q}' . In counterexamples to characterisability, we therefore only provide the smallest natural class of queries within which non-characterisability holds. The following examples illustrate (non-)characterisability within the classes $\mathcal{Q}_p[\diamond_r]$ and $\mathcal{Q}_p[\circ, \diamond_r]$.

Example 1. (i) Recall from Section 1 that $\diamond_r(A \wedge B)$ is not uniquely characterisable within $\mathcal{Q}_p[\diamond_r]$. The same argument shows non-characterisability of $\diamond(A \wedge B)$ within $\mathcal{Q}_p[\diamond_r, \diamond]$. On the other hand, the query $\diamond(A \wedge B)$ is characterised within $\mathcal{Q}_p[\diamond, \circ]$ by the example set with positive examples $(\emptyset, \{A, B\})$ and $(\emptyset, \emptyset, \{A, B\})$ and negative examples $(\emptyset, \{A\})$ and $(\emptyset, \{B\})$.

(ii) The conjunction of atoms does not always lead to non-characterisability within classes of queries with \diamond_r . For example, $q = \diamond_r(A \wedge \circ(A \wedge B))$ is characterised within $\mathcal{Q}_p[\circ, \diamond_r]$ by $E = (E^+, E^-)$ in which E^+ contains two data instances $(\{A\}, \{A, B\})$ and $(\emptyset, \{A\}, \{A, B\})$ and E^- also two instances:

$$(\emptyset, \emptyset, \{A, B\}), \quad (\emptyset, \{A\}, \{A\}, \{B\}, \{A, B\}).$$

The intuition here is that some instances from E^- have to satisfy the query $\diamond_r(A \wedge \circ(B \wedge \diamond_r(A \wedge B)))$ as well as the query $\diamond_r(A \wedge \circ(A \wedge \diamond_r(A \wedge B)))$.

(iii) While the query $\diamond_r(A \wedge B)$ from (i) is not characterisable, there is a polynomial f such that, for all $n \in \mathbb{N}$, it is characterisable within $\mathcal{Q}_p^n[\circ, \diamond_r]$ by some E_n of size $\leq f(n)$. Namely, we take $E^+ = \{(\{A, B\}), (\emptyset, \{A, B\})\}$ and $E^- = \underbrace{\{(\{A\}, \{B\}), \dots, \{A\}, \{B\})\}}_{n \text{ times}}$.

Observe that one can always separate $q \in \mathcal{Q}[\circ, \diamond_r]$ from any other $q' \in \mathcal{Q}[\circ, \diamond_r]$ with $\text{sig}(q') \supsetneq \text{sig}(q) = \sigma$ using the positive example (σ, \dots, σ) with $\text{tdp}(q) + 1$ -many copies of σ . One can therefore focus on characterisability within the relevant class of queries over the same signature as the input query. However, this is not the case for $\mathcal{Q}[\cup]$:

Example 2. The query $q = \perp \cup A \equiv \circ A$ is not uniquely characterisable within $\mathcal{Q}_p[\cup]$. Indeed, suppose q fits E and σ comprises all atoms occurring in E . Then $\mathcal{D}, 0 \models C \cup A$ iff $\mathcal{D}, 0 \models \circ A$, for all \mathcal{D} in E and $C \notin \sigma$, and so E does not characterise q . On the other hand, for the signature $\sigma = \{A, B\}$, the query q is characterised within $\mathcal{Q}_p^\sigma[\cup]$ by the example set (E^+, E^-) in which $E^+ = \{(\emptyset, \{A\})\}$ and $E^- = \{(\sigma, \{B\}, \{A\})\}$ as $A \cup A \equiv (A \wedge B) \cup A \equiv \circ A$.

As noted in Section 1, $\perp U A$ is not uniquely characterisable within $\mathcal{Q}^{\{A\}}[U]$ because of nested U -operators on the left-hand side of U . This observation prompts us to consider the subclass $\mathcal{Q}_-^\sigma[U]$ of $\mathcal{Q}^\sigma[U]$ -queries q in which any subquery $q' \cup q''$ does not contain occurrences of U in q' . Note that $\mathcal{Q}_-^\sigma[U] \subseteq \mathcal{Q}^\sigma[U]$. We show that $\mathcal{Q}_-^\sigma[U]$ is uniquely characterisable. To simplify notation, we give σ -data instances as *words* over the alphabet 2^σ using the standard notation of regular languages. Instead of $\mathcal{D}, 0 \models q$ we simply write $\mathcal{D} \models q$. By the semantics of U , for any $q \in \mathcal{Q}_-^\sigma[U]$, we have

$$\sigma^d \not\models q \text{ for } d \leq \text{tdp}(q), \quad \sigma^d \models q \text{ for } d > \text{tdp}(q) \quad (3)$$

where σ^d is a word with d -many σ . Note also that there are finitely-many, say $N_d < \omega$, pairwise non-equivalent queries of any depth $d < \omega$ in $\mathcal{Q}_-^\sigma(U)$.

Lemma 3. *If $q, q' \in \mathcal{Q}_-^\sigma[U]$ are of depth d and $q \not\models q'$, then there is \mathcal{D} such that $\max(\mathcal{D}) \leq N_d$, $\mathcal{D} \models q$ and $\mathcal{D} \not\models q'$.*

Proof. Consider \mathcal{D} of minimal length such that $\mathcal{D} \models q$ and $\mathcal{D} \not\models q'$. Let $tp(i)$ comprise all of the subqueries s of q and q' with $\mathcal{D}, i \models s$. By the choice of \mathcal{D} , we have $tp(i) \neq tp(j)$ for any distinct $i, j \in [0, \max(\mathcal{D})]$ (otherwise we could cut the interval $[i, j]$ out of \mathcal{D} to obtain a shorter instance separating q from q'). It follows that $\max(\mathcal{D}) \leq N_d$. \square

Theorem 4. *For any σ , $\mathcal{Q}_-^\sigma[U]$ is uniquely characterisable.*

Proof. Any $q \in \mathcal{Q}_-^\sigma(U)$ is uniquely characterised by E with

$$E^+ = \{\mathcal{D} \models q \mid \max(\mathcal{D}) \leq N_{\text{tdp}(q)}\}, \\ E^- = \{\mathcal{D} \not\models q \mid \max(\mathcal{D}) \leq N_{\text{tdp}(q)}\}.$$

Indeed, let $q' \in \mathcal{Q}_-^\sigma(U)$ fit E . Then $\text{tdp}(q') = \text{tdp}(q)$ by (3), and so $q \equiv q'$ by Lemma 3. \square

It follows from the proof that $\mathcal{Q}[\circ, \diamond]$ is uniquely characterisable as well.

5 Characterisability in $\mathcal{Q}_p[\circ, \diamond_r]$

In this section, we prove a criterion of (polynomial) unique characterisability of queries within $\mathcal{Q}_p[\circ, \diamond_r]$. The criterion is applicable to $\mathcal{Q}_p[\circ, \diamond, \diamond_r]$ -queries in a normal form, which is defined and illustrated below.

Example 5. It is readily checked that the $\mathcal{Q}_p[\circ, \diamond_r]$ -query $q = \circ \diamond_r \circ \diamond_r (A \wedge B \wedge C \wedge \diamond_r (B \wedge \diamond_r (B \wedge C)))$ is equivalent to the $\mathcal{Q}_p[\circ, \diamond, \diamond_r]$ -query $q^{nf} = \diamond \diamond (A \wedge B \wedge C)$.

We define the normal form for $\mathcal{Q}_p[\circ, \diamond, \diamond_r]$ -queries represented as a first-order CQ by a list of atoms. For example, the query q^{nf} above is given by the CQ

$$q^{nf}(t_0) = t_0 < t_1, t_1 < t_2, A(t_2), B(t_2), C(t_2)$$

with one free (answer) variable t_0 and existentially quantified t_1 and t_2 . In general, any $q \in \mathcal{Q}_p[\circ, \diamond, \diamond_r]$ is represented as a CQ

$\rho_0(t_0), R_1(t_0, t_1), \dots, \rho_{n-1}(t_{n-1}), R_n(t_{n-1}, t_n), \rho_n(t_n)$, where ρ_i is a set of atoms, $\rho_i(t_i) = \{A(t_i) \mid A \in \rho_i\}$ and $R_i \in \{suc, <, \leq\}$. We divide q into *blocks* q_i such that

$$q = q_0 \mathcal{R}_1 q_1 \dots \mathcal{R}_n q_n \quad (4)$$

with $\mathcal{R}_i = R_1^i(t_0, t_1) \dots R_{n_i}^i(t_{n_i-1}, t_{n_i})$, for $R_j^i \in \{<, \leq\}$,

$$q_i = \rho_0^i(s_0^i) suc(s_0^i, s_1^i) \rho_1^i(s_1^i) \dots suc(s_{k_i-1}^i, s_{k_i}^i) \rho_{k_i}^i(s_{k_i}^i)$$

and $s_{k_i}^i = t_0^{i+1}, t_{n_i}^i = s_0^i$. If $k_i = 0$, the block q_i is *primitive*. A primitive block $q_i = \rho_0^i(s_0^i)$ with $i > 0$ and $|\rho_0^i| \geq 2$ is called a *lone conjunct* of q .

Example 6. The query $\diamond_r(A \wedge B)$ in Example 1(i), whose CQ representation is $t_0 \leq t_1, \rho_1(t_1)$, for $\rho_1 = \{A, B\}$, has a lone conjunct $\rho_1(t_1)$. In $\diamond_r(A \wedge \circ(A \wedge B))$ from Example 1(ii), represented as $t_0 \leq t_1, A(t_1), suc(t_1, t_2), \rho_1(t_2)$, the conjunct $\rho_1(t_2)$ is not lone.

Now, we say that q given by (4) is in *normal form* if the following conditions are satisfied:

- (n1) $\rho_0^i \neq \emptyset$ if $i > 0$, and $\rho_{k_i}^i \neq \emptyset$ if $i > 0$ or $k_i > 0$ (thus, of all the first/last ρ in a block, only ρ_0^0 can be empty);
- (n2) each \mathcal{R}_i is either a single $t_0^i \leq t_1^i$ or a sequence of $<$;
- (n3) $\rho_{k_i}^i \not\supseteq \rho_0^{i+1}$ if q_{i+1} is primitive and R_{i+1} is \leq ;
- (n4) $\rho_{k_i}^i \not\supseteq \rho_0^{i+1}$ if $i > 0$, q_i is primitive and R_{i+1} is \leq .

The queries in Example 6 are in normal form with two blocks each; the query q^{nf} above is in normal form with two blocks $q_0 = \top(t_0)$ and $q_1 = A(t_2) \wedge B(t_2) \wedge C(t_2)$.

Lemma 7. *Every query in $\mathcal{Q}_p[\circ, \diamond_r]$ is equivalent to a query in normal form that can be computed in linear time.*

A query $q \in \mathcal{Q}_p[\circ, \diamond]$ is *safe* if it is equivalent to a query $q' \in \mathcal{Q}_p[\circ, \diamond]$ in normal form not containing lone conjuncts. We are now in the position to formulate the criterion.

Theorem 8. (i) *A query $q \in \mathcal{Q}_p[\circ, \diamond_r]$ is uniquely characterisable within $\mathcal{Q}_p[\circ, \diamond_r]$ iff q is safe.*

(ii) *Those queries that are uniquely characterisable within $\mathcal{Q}_p[\circ, \diamond_r]$ are actually polynomially characterisable within $\mathcal{Q}_p[\circ, \diamond_r]$.*

(iii) *The class $\mathcal{Q}_p[\circ, \diamond_r]$ is polynomially characterisable for bounded query size.*

(iv) *The class $\mathcal{Q}_p[\circ, \diamond]$ is polynomially characterisable.*

Proof sketch. A detailed proof is given in the full version. Here, we define a polysize example set $E = (E^+, E^-)$ characterising a query q in normal form (4), which does not contain lone conjuncts. Let b be the number of \circ and \diamond in q plus 1. For each block q_i in (4), we take two words

$$\bar{q}_i = \rho_0^i \dots \rho_{k_i}^i, \quad \bar{q}_i \bowtie \bar{q}_{i+1} = \rho_0^i \dots (\rho_{k_i}^i \cup \rho_0^{i+1}) \dots \rho_{k_{i+1}}^{i+1}.$$

The set E^+ contains the data instances given by the words

- $\mathcal{D}_b = \bar{q}_0 \emptyset^b \dots \bar{q}_i \emptyset^b \bar{q}_{i+1} \dots \emptyset^b \bar{q}_n$,
- $\mathcal{D}_i = \bar{q}_0 \emptyset^b \dots \bar{q}_i \bowtie \bar{q}_{i+1} \dots \emptyset^b \bar{q}_n$ if \mathcal{R}_{i+1} is \leq ,
- $\mathcal{D}_i = \bar{q}_0 \emptyset^b \dots \bar{q}_i \emptyset^{n_{i+1}} \bar{q}_{i+1} \dots \emptyset^b \bar{q}_n$ otherwise.

Here, \emptyset^b is a sequence of b -many \emptyset and similarly for $\emptyset^{n_{i+1}}$. The set E^- contains all data instances of the form

- $\mathcal{D}_i^- = \bar{q}_0 \emptyset^b \dots \bar{q}_i \emptyset^{n_{i+1}-1} \bar{q}_{i+1} \dots \emptyset^b \bar{q}_n$ if $n_{i+1} > 1$;
- $\mathcal{D}_i^- = \bar{q}_0 \emptyset^b \dots \bar{q}_i \bowtie \bar{q}_{i+1} \dots \emptyset^b \bar{q}_n$ if \mathcal{R}_{i+1} is a single $<$,

and also the data instances obtained from \mathcal{D}_b by

- (a) removing a single atom from some $\rho_j^i \neq \emptyset$ or removing the whole $\rho_j^i = \emptyset$, for $i \neq 0$ and $j \neq 0$, from some \bar{q}_i ;
- (b) replacing $\bar{q}_i = \rho_0^i \dots \rho_l^i \rho_{l+1}^i \dots \rho_{k_i}^i$ ($k_i > 0$) by $\bar{q}'_i \emptyset^b \bar{q}''_i$, where $\bar{q}'_i = \rho_0^i \dots \rho_l^i$, $\bar{q}''_i = \rho_{l+1}^i \dots \rho_{k_i}^i$ and $l \geq 0$;
- (c) replacing some $\rho_l^i \neq \emptyset$, $0 < l < k_i$, by $\rho_l^i \emptyset^b \rho_l^i$;
- (d) replacing $\rho_{k_i}^i$ ($k_i > 0$, $|\rho_{k_i}^i| \geq 2$) with $\rho_{k_i}^i \setminus \{A\} \emptyset^b \rho_{k_i}^i$, for some $A \in \rho_{k_i}^i$, or replacing ρ_0^i ($k_i > 0$, $|\rho_0^i| \geq 2$) with $\rho_0^i \emptyset^b \rho_0^i \setminus \{A\}$, for some $A \in \rho_0^i$;
- (e) replacing $\rho_0^i \neq \emptyset$ with $\rho_0^i \setminus \{A\} \emptyset^b \rho_0^i$, for some $A \in \rho_0^i$, if $k_0 = 0$, and with $\rho_0^i \emptyset^b \rho_0^i$ if $k_0 > 0$.

The size of E is clearly polynomial in $|q|$. It is readily seen that $\mathcal{D} \models q$ for all $\mathcal{D} \in E^+$. To continue the proof sketch, note that $\mathcal{D} \models q$ iff there is a *homomorphism* h from the set $\text{var}(q)$ of variables in q to $[0, \max(\mathcal{D})]$, i.e., $h(t_0) = 0$, $A(h(t)) \in \mathcal{D}$ if $A(t) \in q$, $h(t') = h(t) + 1$ if $\text{suc}(t, t') \in q$, and $h(t) R h(t')$ if $R(t, t') \in q$ for $R \in \{<, \leq\}$. Using the assumption that q is in normal form, one can show that there is no homomorphism witnessing $\mathcal{D} \models q$, for any $\mathcal{D} \in E^-$.


Suppose now that $q' \in \mathcal{Q}_p[\diamond, \diamond_r]$ in normal form is given and $q' \neq q$. If $\mathcal{D}_b \not\models q'$, we are done as $\mathcal{D}_b \in E^+$. Otherwise, let h be a homomorphism witnessing $\mathcal{D}_b \models q'$. Then one can show that either the restriction of h to the blocks of q' is an isomorphism onto the blocks of q or there exists a data instance \mathcal{D} obtained using one of the rules (a)–(e) such that a suitably modified h is a homomorphism from q' to \mathcal{D} . In the latter case, we are done as $\mathcal{D} \in E^-$ and $\mathcal{D} \models q'$. In the former case, q and q' coincide with the exception of the sequences of \diamond and \diamond_r between blocks. Then q can be separated from q' using the examples \mathcal{D}_i and \mathcal{D}_i^- . \square

6 Polynomial Characterisability in $\mathcal{Q}_p^\sigma[\text{U}]$

LTL-queries with U do not correspond to CQs (because of the universal quantification in its semantics), and so require a different approach. We view them as defining regular languages. With each $\mathcal{Q}_p^\sigma[\text{U}]$ -query of the form (2) we associate the following regular expression over the alphabet 2^σ :

$$q = \rho_0 \lambda_1^* \rho_1 \lambda_2^* \dots \lambda_n^* \rho_n \lambda_{n+1}^* \quad (5)$$

where $\lambda_{n+1} = \emptyset$ and $\perp^* = \varepsilon$. We regard the words of the language $L(q)$ over 2^σ as data instances. Clearly, $\mathcal{D}' \models q$ iff there is $\mathcal{D} \in L(q)$ such that $\mathcal{D} \in \mathcal{D}'$, i.e., $\mathcal{D} = (\delta_0, \dots, \delta_k)$ and $\mathcal{D}' = (\delta'_0, \dots, \delta'_k)$, for some $k < \omega$, and $\delta_i \subseteq \delta'_i$, for all $i \leq k$. The language L_q of all σ -data instances $\mathcal{D} \models q$ (regarded as words over 2^σ) can be given by the NFA \mathfrak{A}_q below, where each \rightarrow_α , for $\alpha \neq \perp$, stands for all transitions \rightarrow_β with $\alpha \subseteq \beta \subseteq \sigma$ (note that $\perp \notin \sigma$):



Without loss of generality we assume that all our q are *minimal* in the sense that by replacing any $\lambda_i \neq \perp$ with \perp in q we obtain a query that is *not equivalent* to q . For example, in minimal q , $\rho_j \supseteq \dots \supseteq \rho_i \supseteq \lambda_i$ and $\lambda_l = \perp$ for all $l \in (j, i)$ imply $\rho_j \not\subseteq \lambda_j$ as otherwise $\lambda_j \cup (\rho_j \wedge (\perp \cup \dots \cup (\lambda_i \cup \varphi) \dots))$

is equivalent to $\perp \cup (\rho_j \wedge (\perp \cup \dots \cup (\lambda_i \cup \varphi) \dots))$. Using standard automata-theoretic techniques, one can show:

Theorem 9. Any $\mathcal{Q}_p^\sigma[\text{U}]$ -queries $q \neq q'$ can be separated by some \mathcal{D} with $\max(\mathcal{D}) \leq O((\min\{\text{tdp}(q), \text{tdp}(q')\})^2)$.

Using Theorem 9 in the proof of Theorem 4 we obtain:

Corollary 10. The class $\mathcal{Q}_p^\sigma[\text{U}]$ is exponentially characterisable within $\mathcal{Q}_p^\sigma[\text{U}]$.

The following examples illustrate difficulties in finding short unique characterisations of $\mathcal{Q}_p^\sigma[\text{U}]$ -queries, namely, that in general, data instances of different shapes and forms are needed to separate $\mathcal{Q}_p^\sigma[\text{U}]$ -queries. To unclutter notation we omit $\{\}$ in singletons like $\{A\}$.

Example 11. (a) The shortest data instance separating

$$q = X \emptyset^* A \perp^* B \perp^* A B^* A A^* B \emptyset^*,$$

$$q' = X \emptyset^* A \perp^* B A^* A B^* A \perp^* B \emptyset^*$$

is $\mathcal{D} = X A B A B B A A B$ with $\mathcal{D} \models q$ and $\mathcal{D} \not\models q'$ (e.g., $X A B A B A A B$ satisfies both q and q').

(b) For $l > 0$, let $q_l = (A B^*)^{l-1} A A^* B B^*$. Then

$$X A^* q_{l_1} q_{l_2} \dots q_{l_k} X \emptyset^* \neq X \perp^* q_{l_1} q_{l_2} \dots q_{l_k} X \emptyset^*,$$

$$X A^* q_{l_1} q_{l_2} \dots q_{l_k} A \emptyset^* \equiv X \perp^* q_{l_1} q_{l_2} \dots q_{l_k} A \emptyset^*.$$

If $1 < l_1 \leq \dots \leq l_k$, the former inequivalence is witnessed by the instance $X A^{l_1} B A^{l_2} B \dots A^{l_k} B A^{l_k} B X$. Less generally, $X A^* q_2 q_3 X \emptyset^* \neq X \perp^* q_2 q_3 X \emptyset^*$ can be shown by $X A A B A A A B A A B X$ or by $X A A B A B A A B A B X$ (spot the difference and see (n₂) below).

Here, we consider the class $\mathcal{P}^\sigma[\text{U}]$ of *peerless* queries given by (5), in which, for any i , either $\lambda_i = \perp$ or the sets λ_i and ρ_i are *incomparable* with respect to \subseteq . Our main result is that $\mathcal{P}^\sigma[\text{U}]$ is polynomially characterisable within $\mathcal{Q}_p^\sigma[\text{U}]$.

We start with a general observation. Consider two queries $q = \rho_0 \lambda_1^* \dots \lambda_n^* \rho_n \emptyset^*$ and $q' = \rho_0 \mu_1^* \dots \mu_n^* \rho_n \emptyset^*$. We say that $\lambda_i \neq \perp$ *subsumes* $\mu_j \neq \perp$ if either $i = j$ and $\mu_j \subseteq \lambda_i$, or $j < i$ and $\mu_j \rho_j \dots \rho_{i-1} \subseteq \rho_j \dots \rho_{i-1} \lambda_i$, or $j > i$ and $\rho_i \dots \rho_{j-1} \mu_j \subseteq \lambda_i \rho_i \dots \rho_{j-1}$. In the last two cases,

$$\mu_j \subseteq \rho_j \subseteq \dots \subseteq \rho_{i-1} \subseteq \lambda_i, \quad \mu_j \subseteq \rho_{j-1} \subseteq \dots \subseteq \rho_i \subseteq \lambda_i,$$

respectively. Note that, for peerless q , the last inclusion is impossible. If λ_i and μ_j subsume each other, in which case $\lambda_i = \mu_j$, we call (λ_i, μ_j) a *matching pair*. Observe also that, for $\mathcal{D}_q^i = \rho_0 \dots \rho_{i-1} \lambda_i \rho_i \dots \rho_n$, if $\mathcal{D}_q^i \models q'$, then λ_i subsumes some μ_j : $\rho_0 \dots \rho_n \emptyset \in \mathcal{D}_q^i$ means that λ_i subsumes $\mu_{n+1} = \emptyset$, and $\rho_0 \dots \mu_j \dots \rho_n \in \mathcal{D}_q^i$ that λ_i subsumes μ_j . The proof of the next lemma is given in the full version:

Lemma 12. For any queries q and q' as above, either (i) each $\lambda_i \neq \perp$ subsumes μ_j occurring in some matching pair (λ_k, μ_j) or (ii) q and q' are separated by a data instance of the form \mathcal{D}_q^i or $\mathcal{D}_{q'}^j$. Also, if q is peerless, λ_i can only subsume μ_j in the matching pair (λ_i, μ_j) with $i \geq j$, in which case $\mu_j = \rho_j = \dots = \rho_{i-1} = \lambda_i$.

Note that the number of data instances of the form \mathcal{D}_q^i for all possible $\mathcal{Q}_p^\sigma[\text{U}]$ -queries q' can be exponential in $|\sigma|$. The following example indicates how to overcome this issue.

Example 13. Let $\sigma = \{A, B, C, D, X\}$. To separate the query $X\{C, D\}^*A\emptyset^*$ from any $X\lambda^*A\emptyset^*$ with $A, D \notin \lambda$, we can use $\mathcal{D} = X\sigma \setminus \{A, D\}A$.

Theorem 14. *The class $\mathcal{P}^\sigma[\mathbb{U}]$ is polynomially characterisable within $\mathcal{Q}_p^\sigma[\mathbb{U}]$.*

Proof sketch. We show that any $\mathbf{q} = \rho_0\lambda_1^*\rho_1\lambda_2^*\dots\lambda_n^*\rho_n\emptyset^*$ in $\mathcal{P}^\sigma[\mathbb{U}]$ is characterised by the example set $E = (E^+, E^-)$ where E^+ contains all data instances of the following forms:

- (p₀) $\rho_0 \dots \rho_n$,
- (p₁) $\rho_0 \dots \rho_{i-1}\lambda_i\rho_i \dots \rho_n = \mathcal{D}_q^i$,
- (p₂) $\rho_0 \dots \rho_{i-1}\lambda_i^k\rho_i \dots \rho_{j-1}\lambda_j\rho_j \dots \rho_n = \mathcal{D}_{i,k}^j$, for $i < j$ and $k = 1, 2$;

and E^- has all instances that are *not* in $L(\mathbf{q})$ of the forms:

- (n₀) σ^n and $\sigma^{n-i}\sigma \setminus \{A\}\sigma^i$, for $A \in \rho_i$,
- (n₁) $\rho_0 \dots \rho_{i-1}\sigma \setminus \{A, B\}\rho_i \dots \rho_n$, for $A \in \lambda_i \cup \{\perp\}$ and $B \in \rho_i \cup \{\perp\}$,
- (n₂) for all i and $A \in \lambda_i \cup \{\perp\}$, *some* data instance

$$\mathcal{D}_A^i = \rho_0 \dots \rho_{i-1}(\sigma \setminus \{A\})\rho_i\lambda_{i+1}^{k_{i+1}} \dots \lambda_n^{k_n}\rho_n, \quad (6)$$

if any, such that $\max(\mathcal{D}_A^i) \leq (n+1)^2$ and $\mathcal{D}_A^i \not\models \mathbf{q}^\dagger$ for \mathbf{q}^\dagger obtained from \mathbf{q} by replacing λ_j , for all $j \leq i$, with \perp . Note that $\mathcal{D}_A^i \not\models \mathbf{q}$ for peerless \mathbf{q} .

By definition, \mathbf{q} fits E and $|E|$ is polynomial in $|\mathbf{q}|$. We prove in the full version that E uniquely characterises \mathbf{q} . \square

One reason why this construction does not generalise to the whole $\mathcal{Q}_p^\sigma[\mathbb{U}]$ is that $\mathcal{D}_A^i \not\models \mathbf{q}^\dagger$ does not imply $\mathcal{D}_A^i \not\models \mathbf{q}$ for non-peerless \mathbf{q} , as shown by the following example:

Example 15. Let $\mathbf{q} = XA^*AB^*A\perp^*AB^*AA^*BB^*X\emptyset^*$. For any data instance \mathcal{D}_\perp^3 satisfying (6)—for example, $\mathcal{D}_\perp^3 = XAA\sigma ABABX$ —we have $\mathcal{D}_\perp^3 \models \mathbf{q}$.

7 Characterisability in $\mathcal{Q}[\diamond]$

In the previous two sections, we have investigated characterisability of path-shaped queries. Here, we first justify that restriction by exhibiting two examples that show how temporal branching can destroy polynomial characterisability in $\mathcal{Q}[\diamond]$. Both examples make use of *unbalanced* queries, in which different branches have different length. We then show that this is no accident: one can at least partially restore polynomial characterisability for classes without unbalanced queries.

We start by observing that, without loss of generality, it is enough to consider conjunctions of path queries only:

Lemma 16. *For every $\mathbf{q} \in \mathcal{Q}[\diamond]$, one can compute in polynomial time an equivalent query of the form $\mathbf{q}_1 \wedge \dots \wedge \mathbf{q}_n$ with $\mathbf{q}_i \in \mathcal{Q}_p[\diamond]$, for $1 \leq i \leq n$.*

The first example showing non-polynomial characterisability is rather straightforward but requires unbounded branching and an unbounded number of atoms. We write queries $\mathbf{q} \in \mathcal{Q}_p^\sigma[\diamond]$ of the form

$$\mathbf{q} = \rho_0 \wedge \diamond(\rho_1 \wedge \diamond(\rho_2 \wedge \dots \wedge \diamond\rho_n)) \quad (7)$$

as words $\rho_0\rho_1 \dots \rho_n$ over 2^σ (omitting but not forgetting $\lambda_i^* = \emptyset^*$ from (5)) and also use $\rho_0\rho_1 \dots \rho_n$ to denote the data instance defined by \mathbf{q} .

Example 17. Consider $\mathbf{q}_n = \mathbf{s}_1 \wedge \dots \wedge \mathbf{s}_n$, where $n \geq 2$ and each \mathbf{s}_i is a word repeating n times the sequence $A_1 \dots A_n$ (of singletons) with omitted A_i . Now, consider the queries $\mathbf{q}_n^p = \mathbf{q}_n \wedge \mathbf{p}$, where $\mathbf{p} = \diamond(A_{i_1} \wedge \dots \wedge \diamond(A_{i_2} \wedge \dots \wedge \diamond A_{i_n}))$ and $A_{i_1} \dots A_{i_n}$ is a permutation of $A_1 \dots A_n$. Then $\mathbf{q}_n^p \models \mathbf{q}_n$ and $\mathbf{q}_n \not\models \mathbf{q}_n^p$ as shown by the data instance $\mathbf{s}_{i_1}\mathbf{s}_{i_2} \dots \mathbf{s}_{i_n}$. Moreover, if $\mathcal{D} \models \mathbf{q}_n$, $\mathcal{D} \not\models \mathbf{p}$ and $\mathbf{p}' \neq \mathbf{p}$, then $\mathcal{D} \models \mathbf{p}'$. It follows that, in any $E = (E^+, E^-)$ uniquely characterising \mathbf{q}_n , the set E^+ contains at least $n!$ data instances.

The class $\mathcal{Q}_{\leq n}[\diamond]$ of queries of *branching factor* at most n contains all queries in $\mathcal{Q}[\diamond]$ that are equivalent to a query of the form $\mathbf{q}_1 \wedge \dots \wedge \mathbf{q}_m$ with $m \leq n$ and $\mathbf{q}_i \in \mathcal{Q}_p[\diamond]$. We next provide an example of non-polynomial characterisability that requires four atoms and bounded branching only.

Example 18. Let $\sigma = \{A_1, A_2, B_1, B_2\}$, $\mathbf{q}_1 = \emptyset(\sigma\sigma)^n\mathbf{s}$, and $\mathbf{q}_2 = \emptyset\sigma^{2n+1}$, where $\mathbf{s} = \{A_1, A_2\}\{B_1, B_2\}$. Consider the set P of 2^{n+1} -many queries of the form $\emptyset\mathbf{s}_1 \dots \mathbf{s}_{n+1}$ with \mathbf{s}_i either $\{A_1\}\{A_2\}$ or $\{B_1\}\{B_2\}$. Then $\mathbf{q}_1 \wedge \mathbf{q}_2 \not\models \mathbf{q}$ for any $\mathbf{q} \in P$ and, for any \mathcal{D} with $\mathcal{D} \models \mathbf{q}_1 \wedge \mathbf{q}_2$, there is at most one $\mathbf{q} \in P$ with $\mathcal{D} \models \mathbf{q}$ (the proof is rather involved). It follows that $\mathbf{q}_1 \wedge \mathbf{q}_2 \not\models \mathbf{q}_1 \wedge \mathbf{q}_2 \wedge \mathbf{q}$ for all $\mathbf{q} \in P$, but 2^{n+1} positive examples are needed to separate $\mathbf{q}_1 \wedge \mathbf{q}_2$ from all $\mathbf{q}_1 \wedge \mathbf{q}_2 \wedge \mathbf{q}$ with $\mathbf{q} \in P$.

We next identify polynomially characterisable classes of $\mathcal{Q}[\diamond]$ -queries, assuming as before that $\rho_n \neq \emptyset$ in any \mathbf{q} of the form (1). We call a query $\mathbf{q}_1 \wedge \dots \wedge \mathbf{q}_n \in \mathcal{Q}[\diamond]$ with $\mathbf{q}_1, \dots, \mathbf{q}_n \in \mathcal{Q}_p[\diamond]$ *balanced* if all \mathbf{q}_i have the same depth; further, we call it *simple* if, in each \mathbf{q}_i given by (1), $|\rho_j| = 1$ for all j . Let $\mathcal{Q}_b[\diamond]$ denote the class of queries in $\mathcal{Q}[\diamond]$ that are equivalent to a balanced query.

Theorem 19. (i) *The class of simple queries in $\mathcal{Q}_b[\diamond]$ is polynomially characterisable within $\mathcal{Q}_b[\diamond]$.*

(ii) *For any n , the class $\mathcal{Q}_b[\diamond] \cap \mathcal{Q}_{\leq n}[\diamond]$ is polynomially characterisable.*

Proof sketch. Let $\mathbf{q} \in \mathcal{Q}_p^\sigma[\diamond]$. We start with a lemma on the existence of polynomial-size σ -data instances $\mathcal{D}_{\mathbf{q},k}$ such that $\mathcal{D}_{\mathbf{q},k} \not\models \mathbf{q}$ and $\mathcal{D}_{\mathbf{q},k} \models \mathbf{q}'$ for all $\mathbf{q}' \in \mathcal{Q}_p^\sigma[\diamond]$ with $\mathbf{q}' \not\models \mathbf{q}$ and $\text{tdp}(\mathbf{q}') \leq k$. Note that such $\mathcal{D}_{\mathbf{q},k}$ do not exist in general.

Example 20. Let $\mathbf{q} = A \wedge B$. Then $A \not\models \mathbf{q}$ and $B \not\models \mathbf{q}$ but there does not exist any $\mathcal{D}_{\mathbf{q},0}$ such that $\mathcal{D}_{\mathbf{q},0} \not\models \mathbf{q}$, $\mathcal{D}_{\mathbf{q},0} \models A$ and $\mathcal{D}_{\mathbf{q},0} \models B$.

In the following lemma, we therefore assume that \mathbf{q} does not speak about the initial timepoint.

Lemma 21. *Let $\mathbf{q} \in \mathcal{Q}_p^\sigma[\diamond]$ be of the form $\diamond\mathbf{q}'$ and let $k > 0$. Then one can construct in polynomial time a σ -data instance $\mathcal{D}_{\mathbf{q},k}$ such that $\mathcal{D}_{\mathbf{q},k} \not\models \mathbf{q}$ and $\mathcal{D}_{\mathbf{q},k} \models \mathbf{q}'$ for all $\mathbf{q}' \in \mathcal{Q}_p^\sigma[\diamond]$ with $\mathbf{q}' \not\models \mathbf{q}$ and $\text{tdp}(\mathbf{q}') \leq k$.*

Proof. Assuming that $\mathbf{q} = \diamond(\rho_1 \wedge \diamond(\rho_2 \wedge \dots \wedge \diamond\rho_n))$ with $\rho_i = \{A_1^i, \dots, A_{n_i}^i\}$ for $i \geq 1$, we set

$$\mathcal{D}_{\mathbf{q},k} = \sigma\mathbf{s}_1^k\sigma \dots \mathbf{s}_{n-1}^k\sigma\mathbf{s}_n^k,$$

where $s_i = \sigma \setminus \{A_1^i\} \dots \sigma \setminus \{A_{n_i}^i\}$. One can show by induction that $\mathcal{D}_{q,k}$ is as required. \square

Using Lemma 21, for any $q \in \mathcal{Q}^\sigma[\diamond]$, one can construct a polynomial-size set of negative examples as follows. Suppose $q = q_1 \wedge \dots \wedge q_n \in \mathcal{Q}^\sigma[\diamond]$ with

$$q_i = \rho_0^i \wedge \diamond(\rho_1^i \wedge \diamond(\rho_2^i \wedge \dots \wedge \diamond \rho_{n_i}^i)).$$

Let $\rho = \bigwedge_{i=1}^n \rho_0^i$ and let q_i^- be q_i without the conjunct ρ_0^i , so Lemma 21 is applicable to q_i^- . Now let $E_{q_i, m}^-$ contain the σ -data instances $\mathcal{D}_{q_i^-, m}$ and $\sigma \setminus \{A\} \sigma^m$ for all $A \in \rho$.

Lemma 22. (i) For any $\mathcal{D} \in E_{q, m}^-$, we have $\mathcal{D} \not\models q$.

(ii) For any $q' \in \mathcal{Q}^\sigma[\diamond]$ with $q' \not\models q$ and $\text{tdp}(q') \leq m$, there exists $\mathcal{D} \in E_{q, m}^-$ with $\mathcal{D} \models q'$.

It follows from Lemma 22 that non-polynomial characterisability of $\mathcal{Q}[\diamond]$ -queries can only be caused by the need for super-polynomially-many positive examples. We now discuss the construction of positive examples in the proof of Theorem 19 (ii); part (i) is dealt with in the full version. Let $q = q_1 \wedge \dots \wedge q_m \in \mathcal{Q}_b^\sigma[\diamond] \cap \mathcal{Q}_{\leq n}^\sigma[\diamond]$ with $m \leq n$ and

$$q_i = \rho_0^i \wedge \diamond(\rho_1^i \wedge \diamond(\rho_2^i \wedge \dots \wedge \diamond \rho_N^i)).$$

For any map $f: \{1, \dots, m\} \rightarrow \{1, \dots, N\}$, construct a σ -data instance \mathcal{D}_f by inserting $\rho_{f(i)}^i$ into the data instance σ^N in position $f(i)$. Let E^+ contain the data instance $\rho \sigma^N$ for $\rho = \bigcup_{i=1}^m \rho_0^i$ and all the data instances \mathcal{D}_f . One can show that (E^+, E^-) characterises q in $\mathcal{Q}_b[\diamond] \cap \mathcal{Q}_{\leq n}[\diamond]$. \square

8 2D Temporal Instance Queries

Now we consider ‘two-dimensional’ query languages that combine instance queries (over the object domain) in the standard description logics \mathcal{EL} and \mathcal{ELI} (Baader et al. 2017) with the *LTL*-queries (over the temporal domain) considered above. Our aim is to understand how far the characterisability results of the previous sections can be generalised to the 2D case. A *relational signature* is a finite set $\Sigma \neq \emptyset$ of unary and binary predicate symbols. A Σ -*data instance* \mathcal{A} is a finite set of *atoms* $A(a)$ and $P(a, b)$ with $A, P \in \Sigma$ and *individual names* a, b . Let $\text{ind}(\mathcal{A})$ be the set of individual names in \mathcal{A} . We assume that $P^-(a, b) \in \mathcal{A}$ iff $P(b, a) \in \mathcal{A}$, calling P^- the *inverse* of P (with $P^{- -} = P$). Let $S := P \mid P^-$. *Temporal instance queries* are defined by the grammar

$$q := \top \mid \perp \mid A \mid \exists S.q \mid q_1 \wedge q_2 \mid \text{op } q \mid q_1 \cup q_2,$$

where $\text{op} \in \{\circ, \diamond, \diamond_r\}$. Such queries without temporal operators are called *ELI-queries*; those of them without inverses P^- are *EL-queries*. A *temporal Σ -data instance* \mathcal{D} is a finite sequence $\mathcal{A}_0, \dots, \mathcal{A}_n$ of Σ -data instances. We set $\text{ind}(\mathcal{D}) = \bigcup_{i=1}^n \text{ind}(\mathcal{A}_i)$. For any $\ell \in \mathbb{N}$ and $a \in \text{ind}(\mathcal{D})$, the *truth-relation* $\mathcal{D}, a, \ell \models q$ is defined by induction:

$$\mathcal{D}, a, \ell \models A \text{ iff } A(a) \in \mathcal{A}_\ell,$$

$$\mathcal{D}, a, \ell \models \exists S.q \text{ iff there is } b \in \text{ind}(\mathcal{A}_\ell) \text{ such that}$$

$$S(a, b) \in \mathcal{A}_\ell \text{ and } \mathcal{D}, b, \ell \models q,$$

with the remaining clauses being obvious generalisations of the *LTL* ones. An *example set* is a pair $E = (E^+, E^-)$

with finite sets E^+ and E^- of *pointed temporal data instances* (\mathcal{D}, a) such that $a \in \text{ind}(\mathcal{D})$. We say that q fits E if $\mathcal{D}^+, a^+, 0 \models q$ and $\mathcal{D}^-, a^-, 0 \not\models q$, for all $(\mathcal{D}^+, a^+) \in E^+$ and $(\mathcal{D}^-, a^-) \in E^-$. As before, E *uniquely characterises* q if q fits it and every q' fitting E is logically equivalent to q .

We need the following result on the unique characterisability of \mathcal{ELI} -queries.

Theorem 23 (ten Cate and Dalmau 2021). *The class of \mathcal{ELI} -queries is polynomially characterisable.*

Theorem 24 is proved by constructing frontiers in the set of \mathcal{ELI} -queries partially ordered by entailment, where a set \mathcal{F} of \mathcal{ELI} -queries is called a *frontier* of an \mathcal{ELI} -query q if the following hold:

- $q \models q'$ and $q' \not\models q$, for all $q' \in \mathcal{F}$;
- if $q \models q''$ for some \mathcal{ELI} -query q'' , then $q'' \models q$ or there exists $q' \in \mathcal{F}$ with $q' \models q''$.

Theorem 24 (ten Cate and Dalmau 2021). *A frontier $\mathcal{F}(q)$ of any \mathcal{ELI} -query q can be computed in polynomial time.*

Theorem 23 follows from Theorem 24. For any \mathcal{ELI} -query q , we denote by \hat{q} the tree-shaped data instance defined by q with designated root a . Then q is characterised by E with $E^+ = \{\hat{q}\}$ and $E^- = \{\hat{r} \mid r \in \mathcal{F}(q)\}$.

For any unrestricted temporal query language $\mathcal{Q}[\Phi]$ and $\mathcal{L} \in \{\mathcal{EL}, \mathcal{ELI}\}$, we denote by $\mathcal{Q}[\Phi] \otimes \mathcal{L}$ the set of all temporal instance queries with operators in Φ with (for \mathcal{ELI}) or without (for \mathcal{EL}) inverse predicates. We generalise the path-shaped queries $\mathcal{Q}_p[\Phi]$ as follows: $\mathcal{Q}_p[\Phi] \otimes \mathcal{L}$ denotes the class of queries q in $\mathcal{Q}[\Phi] \otimes \mathcal{L}$ such that, for any subquery $q_1 \wedge q_2$ of q , either q_1 or q_2 do not have an occurrence of any operator in Φ that is not in the scope of $\exists S$. To illustrate, $\exists S.\diamond A \wedge \diamond \exists S.A$ is in $\mathcal{Q}_p[\Phi] \otimes \mathcal{L}$, but $\diamond A \wedge \diamond \exists S.A$ is not. We make two observations about unique characterisability in these ‘full’ combinations.

Theorem 25. (i) $\mathcal{Q}[\circ, \diamond] \otimes \mathcal{EL}$ is uniquely characterisable.

(ii) $\mathcal{Q}_p[\circ] \otimes \mathcal{ELI}$ and $\mathcal{Q}_p[\diamond] \otimes \mathcal{ELI}$ are polynomially characterisable.

Here, (i) is shown similarly to Theorem 4 (it remains open whether it can be extended to $\mathcal{Q}[\circ, \diamond] \otimes \mathcal{ELI}$); (ii) is proved by generalising Theorem 24 to temporal data instances.

We now show that the application of the DL constructor $\exists P$ to temporal queries with both \circ and \diamond destroys polynomial characterisability. Denote by $\mathcal{EL}(\mathcal{Q}_p[\circ, \diamond])$ the class of queries in $\mathcal{Q}_p[\circ, \diamond] \otimes \mathcal{EL}$ that contain no $\exists P$ in the scope of a temporal operator.

Theorem 26. $\mathcal{EL}(\mathcal{Q}_p[\circ, \diamond])$ is not polynomially characterisable.

Proof sketch. Consider queries $q_n = \exists P.q_1^n \wedge \dots \wedge \exists P.q_n^n$, in which each q_i^n corresponds to the regular expression

$$\underbrace{BB\emptyset^*A}_1 \dots \underbrace{BB\emptyset^*A\emptyset^*B\emptyset^*A}_{i-1} \underbrace{\emptyset^*B\emptyset^*A}_{i} \underbrace{BB\emptyset^*A}_{i+1} \dots \underbrace{BB\emptyset^*A\emptyset^*}_{n}$$

(with omitted $\perp^* = \varepsilon$ in BB). One can show that any unique characterisation of q_n contains at least 2^n positive examples to separate it from all queries $q_n \wedge \exists P.s$ with

$$s = \mathbf{o}_1(B \wedge \diamond(A \wedge \mathbf{o}_2(B \wedge \diamond(A \wedge \dots \wedge \mathbf{o}_n(B \wedge \diamond A) \dots))),$$

where \mathbf{o}_i is \circ or $\diamond \circ$ if $i > 1$, and blank or \diamond if $i = 1$. \square

The situation changes drastically if we do not admit temporal operators in the scope of $\exists P$. We start by investigating the class $\mathcal{Q}_p[\circ, \diamond_r](\mathcal{ELI})$ of queries of the form

$$\mathbf{q} = \mathbf{r}_0 \wedge \mathbf{o}_1(\mathbf{r}_1 \wedge \mathbf{o}_2(\mathbf{r}_2 \wedge \dots \wedge \mathbf{o}_n \mathbf{r}_n)),$$

where the \mathbf{r}_i are \mathcal{ELI} -queries and $\mathbf{o}_i \in \{\circ, \diamond_r\}$. We can generalise the CQ-representation, the normal form, and the notion of lone conjunct from $\mathcal{Q}_p[\circ, \diamond_r]$ to $\mathcal{Q}_p[\circ, \diamond_r](\mathcal{ELI})$ in a straightforward way. To formulate conditions (n1)–(n4), we replace the set inclusions ' $\rho_i \subseteq \rho_j$ ' by entailment ' $\mathbf{r}_i \models \mathbf{r}_j$ '. For example, (n4) becomes

$$(n4') \mathbf{r}_0^{i+1} \not\models \mathbf{r}_{k_i}^i \text{ if } i > 0, \mathbf{q}_i \text{ is primitive and } R_{i+1} \text{ is } \leq.$$

The condition for lone conjuncts now requires that \mathbf{r} is not equivalent to any $\mathbf{q}_1 \wedge \mathbf{q}_2$ with \mathcal{ELI} -queries $\mathbf{q}_1, \mathbf{q}_2$ such that $\mathbf{q}_i \not\models \mathbf{r}$ for $i = 1, 2$. Then one can show again that every $\mathcal{Q}_p[\circ, \diamond_r](\mathcal{ELI})$ -query is equivalent to a query in normal form, which can be computed in polynomial time.

Theorem 27. *The statements of Theorem 8 (i)–(iv) also hold if one replaces $\mathcal{Q}_p[\circ, \diamond_r]$ by $\mathcal{Q}_p[\circ, \diamond_r](\mathcal{ELI})$.*

The proof generalises the example set defined in Theorem 8 using the frontiers provided by Theorem 24 as a *black box*. Indeed, in the definition of examples, replace ρ_i by $\hat{\mathbf{r}}_i$, the data instance corresponding to the \mathcal{ELI} -query \mathbf{r}_i , and replace ' $\rho \setminus \{A\}$ for $A \in \rho$ ' by 'the data instance corresponding to a query in $\mathcal{F}(\mathbf{r})$ '. We choose a single individual, a , as the root of these data instances. For example, item (a) becomes:

(a') replacing some \mathbf{r}_j^i by the data instance corresponding to a query in $\mathcal{F}(\mathbf{r}_j^i)$ or removing the whole $\mathbf{r}_j^i = \emptyset$ for $i \neq 0$ and $j \neq 0$ from some \mathbf{q}_i .

Next, consider the class $\mathcal{Q}_p[\cup](\mathcal{L})$ of queries of the form

$$\mathbf{q} = \mathbf{r}_0 \wedge (\mathbf{l}_1 \cup (\mathbf{r}_1 \wedge (\mathbf{l}_2 \cup (\dots (\mathbf{l}_n \cup \mathbf{r}_n) \dots))), \quad (8)$$

where \mathbf{r}_i is an \mathcal{L} -query and \mathbf{l}_i is either an \mathcal{L} -query or \perp , for $\mathcal{L} \in \{\mathcal{EL}, \mathcal{ELI}\}$. For the same reason as in the 1D case, we fix a finite signature Σ of predicate symbols. Denote by $\mathcal{L}(\Sigma)$ and $\mathcal{Q}_p[\cup](\mathcal{L})$ the set of queries in \mathcal{L} and $\mathcal{Q}_p^\Sigma[\cup](\mathcal{L})$, respectively, with predicate symbols in Σ . Aiming to generalise Theorem 14, we again translate set-inclusion to entailment, so the *peerless queries* $\mathcal{P}^\Sigma[\cup](\mathcal{L})$ take the form (8) such that either $\mathbf{l}_i = \perp$ or $\mathbf{l}_i \not\models \mathbf{r}_i$ and $\mathbf{r}_i \not\models \mathbf{l}_i$.

Theorem 28. *Let Σ be a finite relational signature. Then $\mathcal{P}^\Sigma[\cup](\mathcal{EL})$ is polynomially characterisable within $\mathcal{Q}_p^\Sigma[\cup](\mathcal{EL})$, while $\mathcal{P}^\Sigma[\cup](\mathcal{ELI})$ is exponentially, but not polynomially, characterisable within $\mathcal{Q}_p^\Sigma[\cup](\mathcal{ELI})$.*

To prove Theorem 28, we generalise the example set from the proof of Theorem 14. The positive examples are straightforward: simply replace ρ_i and λ_i by the data instances corresponding to \mathbf{r}_i and \mathbf{l}_i (and choose a single root individual). For the negative examples, we have to generalise the construction of σ , $\sigma \setminus \{A\}$, and $\sigma \setminus \{A, B\}$. For σ , this is straightforward as its role can now be played by the Σ -data instance $\mathcal{A}_\Sigma = \{A(a), R(a, a) \mid A, R \in \Sigma\}$ for which $\mathcal{A}_\Sigma \models \mathbf{q}(a)$ for all $\mathbf{q} \in \mathcal{ELI}(\Sigma)$. For $\sigma \setminus \{A\}$ and $\sigma \setminus \{A, B\}$, we require *split partners* defined as follows. Let Q be a finite set of $\mathcal{L}(\Sigma)$ -queries. A set $\mathcal{S}(Q)$ of pointed Σ -data instances (\mathcal{A}, a) is called a *split partner* of Q in $\mathcal{L}(\Sigma)$ if the following conditions are equivalent for all $\mathcal{L}(\Sigma)$ -queries \mathbf{q}' :

- $\mathcal{A} \models \mathbf{q}'(a)$ for some $(\mathcal{A}, a) \in \mathcal{S}(Q)$;
- $\mathbf{q}' \not\models \mathbf{q}$ for all $\mathbf{q} \in Q$.

Example 29. The split partner of $\{\mathbf{q} = A\}$ in $\mathcal{EL}(\Sigma)$ is the singleton set containing $(\mathcal{A}_\Sigma^{-A}, a)$ with \mathcal{A}_Σ^{-A} defined as $\{B(a), R(a, b), R(b, b), B'(b) \mid B \in \Sigma \setminus \{A\}, R, B' \in \Sigma\}$.

Theorem 30. *Fix $n > 0$. For any set Q of $\mathcal{EL}(\Sigma)$ -queries with $|Q| \leq n$, one can compute in polynomial time a split partner $\mathcal{S}(Q)$ of Q in $\mathcal{EL}(\Sigma)$. For \mathcal{ELI} , one can compute a split partner in exponential time, which is optimal as even for singleton sets Q of $\mathcal{ELI}(\Sigma)$ -queries, no polynomial-size split partner of Q in $\mathcal{ELI}(\Sigma)$ exists in general.*

The proof, given in the full version, requires (as does \mathcal{A}_Σ) the construction of non-tree-shaped data instances. Our results for \mathcal{ELI} are closely related to the study of generalised dualities for homomorphisms between relational structures (Foniok, Nesetril, and Tardif 2008; Nesetril and Tardif 2005) but use pointed relational structures. The construction of $\mathcal{S}(Q)$ for \mathcal{ELI} is based on a construction first introduced in (Bienvenu et al. 2014).

We obtain the negative examples for \mathbf{q} of the form (8) by taking the following pointed data instances (\mathcal{D}, a) (assuming that split partners take the form (\mathcal{A}, a) for a fixed a):

- (n₀') $(\mathcal{A}_\Sigma^n, a)$ and $(\mathcal{A}_\Sigma^{n-i} \mathcal{A} \mathcal{A}_\Sigma^i, a)$, for $(\mathcal{A}, a) \in \mathcal{S}(\{\mathbf{r}_i\})$;
- (n₁') $(\mathcal{D}, a) = (\hat{\mathbf{r}}_0 \dots \hat{\mathbf{r}}_{i-1} \mathcal{A} \hat{\mathbf{r}}_i \dots \hat{\mathbf{r}}_n, a)$ with $\mathcal{D}, a, 0 \not\models \mathbf{q}$ and $(\mathcal{A}, a) \in \mathcal{S}(\{\mathbf{l}_i, \mathbf{r}_i\}) \cup \mathcal{S}(\{\mathbf{l}_i\}) \cup \mathcal{S}(\{\mathbf{r}_i\}) \cup \{(\mathcal{A}_\Sigma, a)\}$;
- (n₂') for all i and $(\mathcal{A}, a) \in \mathcal{S}(\{\mathbf{l}_i\}) \cup \{(\mathcal{A}_\Sigma, a)\}$, some data instance

$$(\mathcal{D}_{\mathcal{A}}^i, a) = (\hat{\mathbf{r}}_0 \dots \hat{\mathbf{r}}_{i-1} \mathcal{A} \hat{\mathbf{r}}_i \hat{\mathbf{l}}_{i+1}^{k_{i+1}} \hat{\mathbf{r}}_{i+1} \dots \hat{\mathbf{l}}_n^{k_n} \mathbf{r}_n, a),$$

if any, such that $\max(\mathcal{D}_{\mathcal{A}}^i) \leq (n+1)^2$ and $\mathcal{D}_{\mathcal{A}}^i, a, 0 \not\models \mathbf{q}^\dagger$ for \mathbf{q}^\dagger obtained from \mathbf{q} by replacing all \mathbf{l}_j , $j \leq i$, with \perp .

We illustrate the construction by generalising Example 2.

Example 31. For $\mathbf{q} = \circ A$ and any relational signature $\Sigma \ni A$, we obtain, after removing redundant instances, that $E^+ = \{(\emptyset\{A(a)\}, a)\}$ and $E^- = \{(\mathcal{A}_\Sigma \mathcal{A}_\Sigma^{-A}\{A(a)\}, a)\}$ characterise \mathbf{q} within $\mathcal{Q}_p^\Sigma[\cup](\mathcal{EL})$.

We finally generalise Theorem 19 (ii) (part (i) is not interesting since simple queries do not generalise to any new class of \mathcal{ELI} -queries). Query classes such as $\mathcal{Q}[\diamond](\mathcal{EL})$ are defined in the obvious way by replacing in $\mathcal{Q}[\diamond]$ -queries conjunctions of atoms by \mathcal{EL} -queries.

Theorem 32. *The class $\mathcal{Q}_b[\diamond](\mathcal{EL}) \cap \mathcal{Q}_{\leq n}[\diamond](\mathcal{EL})$ is polynomially characterisable for any $n < \omega$.*

Again the positive and negative examples are obtained from the 1D case by replacing σ by \mathcal{A}_Σ and $\sigma \setminus \{A\}$ by appropriate split partners.

9 Applications to Learning

We apply our results on unique characterisability to exact learnability of temporal instance queries. Given a class \mathcal{Q} of such queries, we aim to identify a *target query* $\mathbf{q} \in \mathcal{Q}$ using queries to an oracle. The learner knows \mathcal{Q} and the signature σ (Σ in the 2D case) of \mathbf{q} . We allow only one type of queries,

called *membership queries*, in which the learner picks a σ -data instance \mathcal{D} and asks the oracle whether $\mathcal{D} \models \mathbf{q}$ holds. (In the 2D case, the learner picks a pointed Σ -data instance (\mathcal{D}, a) and asks whether $\mathcal{D}, a, 0 \models \mathbf{q}$ holds.) The oracle answers ‘yes’ or ‘no’ truthfully. The class \mathcal{Q} is (*polynomial time*) *learnable with membership queries* if there exists an algorithm that halts for any $\mathbf{q} \in \mathcal{Q}$ and computes (in polynomial time in the size of \mathbf{q} and σ/Σ), using membership queries, a query $\mathbf{q}' \in \mathcal{Q}$ that is equivalent to \mathbf{q} . By default, the learner does not know $|\mathbf{q}|$ in advance but reflecting Theorem 8 (iii), we also consider the case when $|\mathbf{q}|$ is known (which is common in active learning).

Obviously, unique characterisability is a necessary condition for learnability with membership queries. Conversely, if there is an algorithm that computes, for every $\mathbf{q} \in \mathcal{Q}$, an example set that uniquely characterises \mathbf{q} within $\mathcal{Q}^{sig(\mathbf{q})}$, then \mathcal{Q} is learnable with membership queries: enumerate $\mathcal{Q}^{sig(\mathbf{q})}$ starting with the smallest query \mathbf{q} , compute a characterising set E for \mathbf{q} and check using membership queries whether \mathbf{q} is equivalent to the target query. Eventually the algorithm will terminate with a query that is equivalent to the target query. As all of our positive results on unique characterisability provide algorithms computing example sets, we directly obtain learnability with membership queries. Moreover, if the example sets are computed in exponential time, then we obtain an exponential-time learning algorithm: in the enumeration above only $|sig(\mathbf{q})|^{|\mathbf{q}|}$ queries are checked before the target query is found. Unfortunately, we cannot infer polynomial-time learnability from polynomial characterisability in this way.

A detailed analysis of polynomial-time learnability using membership queries is beyond the scope of this paper. Instead, we focus on one main result, the polynomial-time learnability of $\mathcal{Q}_p[\circ, \diamond](\mathcal{ELT})$.

Theorem 33. (i) *The class of safe queries in $\mathcal{Q}_p[\circ, \diamond_r](\mathcal{ELT})$ is polynomial-time learnable with membership queries.*

(ii) *The class $\mathcal{Q}_p[\circ, \diamond_r](\mathcal{ELT})$ is polynomial-time learnable with membership queries if the learner knows the size of the target query in advance.*

(iii) *The class $\mathcal{Q}_p[\circ, \diamond](\mathcal{ELT})$ is polynomially-time learnable with membership queries.*

Proof sketch. We consider the 1D case without \mathcal{ELT} -queries first. (i) Our proof strategy is to construct a query \mathbf{q}' that agrees with \mathbf{q} on the positive and negative examples for \mathbf{q}' from Theorem 8. The algorithm proceeds by computing a data instance \mathcal{D} . Our aim is to arrive at \mathcal{D}_b through iterations of steps, from which the required query can be ‘read off’.

Step 1. First, identify the number of \circ and \diamond in \mathbf{q} by asking membership queries of the form σ^k incrementally, starting from $k = 1$, and then set $b = \min\{k \mid \sigma^k \models \mathbf{q}\} + 1$ and $\mathcal{D}_0 = \sigma^b$. Initialise $\mathcal{D} = \mathcal{D}_0$.

Step 2. Suppose that a data instance \mathcal{D}' is obtained from \mathcal{D} by applying one of the rules (a)–(e) of Theorem 8. If $\mathcal{D}' \models \mathbf{q}$ then replace \mathcal{D} with \mathcal{D}' . Repeat as long as possible. One can show that the number of applications

of each rule is bounded by a polynomial in $|\sigma|$ and the size of \mathbf{q} , and so **Step 2** finishes in polynomial time.

Step 3. Suppose \mathcal{D} contains $\emptyset^b \rho_0^i \emptyset^b$ and $|\rho_0^i| \geq 2$. Since rule (a) does not apply, every homomorphism $h: \mathbf{q} \rightarrow \mathcal{D}$ sends some t_1, \dots, t_l to ρ_0^i , for $l \geq 1$. As \mathbf{q} does not contain lone conjuncts, \mathbf{q} contains singleton primitive blocks at positions t_1, \dots, t_l . Suppose $\rho_0^i = \{A_1, \dots, A_{|\rho_0^i|}\}$ and let $w = \{A_1\} \emptyset^b \{A_2\} \emptyset^b \dots \{A_{|\rho_0^i|}\} \emptyset^b$ (the order in which $A_1, \dots, A_{|\rho_0^i|}$, the elements of ρ_0^i , are enumerated does not matter, we fix any one). Let \mathcal{D}_k^i be obtained from \mathcal{D} by replacing $\emptyset^b \rho_0^i \emptyset^b$ with $\emptyset^b(w)^k$. Notice that, for $k = |\mathbf{q}|$, we have $\mathcal{D}_k^i \models \mathbf{q}$; however, the algorithm is not given this k . Instead, the algorithm incrementally iterates starting from $k = 1$ until $\mathcal{D}_k^i \models \mathbf{q}$. Since $k \leq |\mathbf{q}|$, this takes polynomially-many iterations. Let \mathcal{D}' be obtained from \mathcal{D}_k^i by removing primitive blocks as long as $\mathcal{D}' \models \mathbf{q}$. Notice that rules (a)–(e) do not apply to \mathcal{D}' . Replace \mathcal{D} with \mathcal{D}' . Repeat **Step 3** as long as possible. Since no new lone conjuncts are introduced, the process finishes after polynomially-many steps.

Step 4. At this point of computation, the algorithm has identified all blocks of \mathbf{q} but not the sequences of \diamond and \diamond_r between them. They can be easily determined based on the positive and negative examples \mathcal{D}_i and \mathcal{D}_i^- .

The proof of (ii) is similar, with a modified **Step 3**. Finally, (iii) is a consequence of (ii) as the size of the query \mathbf{q} does not exceed $n = |\sigma|b$.

We obtain a learning algorithm for $\mathcal{Q}_p[\circ, \diamond_r](\mathcal{ELT})$ by combining the learning algorithm above with the learning algorithm for \mathcal{ELT} -queries by ten Cate and Dalmau (2021) using the positive and negative examples given in Theorem 27. Note that the data instance \mathcal{A}_Σ is now used instead of σ and that one has to ‘unfold’ such non tree-shaped data instances into tree-shaped ones. \square

10 Conclusions

In this paper, we have considered temporal instance queries with *LTL* operators and started investigating their unique (polynomial) characterisability and exact learnability using membership queries. We have obtained both positive and negative results, depending on the available temporal operators and the allowed interaction between the temporal and object dimensions in queries. The results indicate that finding complete classifications of 1D and 2D temporal queries according to (polynomial) characterisability and learnability could be a very difficult task. In particular, interesting open problems include the polynomial characterisability of full $\mathcal{Q}_p^\sigma[\cup]$, more general criteria of polynomial characterisability for temporal branching queries and other temporal operators, and the polynomial-time learnability of $\mathcal{Q}_p^\sigma[\cup]$ and 2D extensions. From a conceptual viewpoint, it would be of interest to develop a framework that spells out explicitly the conditions that non-temporal queries should satisfy so that their combination with *LTL*-queries preserves polynomial characterisability and polynomial-time learnability.

Acknowledgments

This research was supported by the EPSRC UK grants EP/S032207 and EP/S032282 for the joint project ‘quant^{MD}: Ontology-Based Management for Many-Dimensional Quantitative Data’.

References

- Aarts, F., and Vaandrager, F. 2010. Learning i/o automata. In *International Conference on Concurrency Theory*, 71–85. Springer.
- Alexe, B.; ten Cate, B.; Kolaitis, P. G.; and Tan, W. C. 2011. Characterizing schema mappings via data examples. *ACM Trans. Database Syst.* 36(4):23.
- Angluin, D.; Frazier, M.; and Pitt, L. 1992. Learning conjunctions of Horn clauses. *Mach. Learn.* 9:147–164.
- Angluin, D. 1987a. Learning regular sets from queries and counterexamples. *Inf. Comput.* 75(2):87–106.
- Angluin, D. 1987b. Queries and concept learning. *Mach. Learn.* 2(4):319–342.
- Arenas, M., and Diaz, G. I. 2016. The exact complexity of the first-order logic definability problem. *ACM Trans. Database Syst.* 41(2):13:1–13:14.
- Artale, A.; Kontchakov, R.; Kovtunova, A.; Ryzhikov, V.; Wolter, F.; and Zakharyashev, M. 2017. Ontology-mediated query answering over temporal data: A survey. In *Proc. of TIME 2017*, volume 90 of *LIPICs*, 1:1–1:37. Schloss Dagstuhl, Leibniz-Zentrum für Informatik.
- Baader, F.; Horrocks, I.; Lutz, C.; and Sattler, U. 2017. *An Introduction to Description Logics*. Cambridge University Press.
- Baader, F.; Borgwardt, S.; and Lippmann, M. 2015. Temporal query entailment in the description logic SHQ. *J. Web Semant.* 33:71–93.
- Barceló, P., and Romero, M. 2017. The complexity of reverse engineering problems for conjunctive queries. In *Proc. of ICDT*, 7:1–7:17.
- Bienvenu, M.; ten Cate, B.; Lutz, C.; and Wolter, F. 2014. Ontology-based data access: A study through disjunctive datalog, CSP, and MMSNP. *ACM Trans. Database Syst.* 39(4):33:1–33:44.
- Borgwardt, S., and Thost, V. 2015. Temporal query answering in the description logic EL. In *Proc. of IJCAI 2015*, 2819–2825. AAAI Press.
- Camacho, A., and McIlraith, S. A. 2019. Learning interpretable models expressed in linear temporal logic. In *Proc. of ICAPS 2018*, 621–630. AAAI Press.
- Cassel, S.; Howar, F.; Jonsson, B.; and Steffen, B. 2016. Active learning for extended finite state machines. *Formal Aspects Comput.* 28(2):233–263.
- Chomicki, J., and Toman, D. 2018. *Temporal Logic in Database Query Languages*. Springer. 3992–3998.
- Fijalkow, N., and Lagarde, G. 2021. The complexity of learning linear temporal formulas from examples. *CoRR* abs/2102.00876.
- Foniok, J.; Nesetril, J.; and Tardif, C. 2008. Generalised dualities and maximal finite antichains in the homomorphism order of relational structures. *Eur. J. Comb.* 29(4):881–899.
- Fortin, M.; Konev, B.; Ryzhikov, V.; Savateev, Y.; Wolter, F.; and Zakharyashev, M. 2022a. Reverse engineering of temporal queries with and without LTL ontologies: First steps (extended abstract). In *Proc. of DL’22*.
- Fortin, M.; Konev, B.; Ryzhikov, V.; Savateev, Y.; Wolter, F.; and Zakharyashev, M. 2022b. Unique Characterisability and Learnability of Temporal Instance Queries. *CoRR* abs/2205.01651.
- Funk, M.; Jung, J. C.; Lutz, C.; Pulcini, H.; and Wolter, F. 2019. Learning description logic concepts: When can positive and negative examples be separated? In *Proc. of IJCAI*, 1682–1688.
- Funk, M.; Jung, J. C.; and Lutz, C. 2021. Actively learning concepts and conjunctive queries under ELr-ontologies. In *Proc. of IJCAI 2021*, 1887–1893. ijcai.org.
- Gutiérrez-Basulto, V.; Jung, J. C.; and Sabellek, L. 2018. Reverse engineering queries in ontology-enriched systems: The case of expressive Horn description logic ontologies. In *Proc. of IJCAI-ECAI*.
- Hodkinson, I. M.; Wolter, F.; and Zakharyashev, M. 2000. Decidable fragment of first-order temporal logics. *Ann. Pure Appl. Log.* 106(1-3):85–134.
- Howar, F., and Steffen, B. 2018. Active automata learning in practice - an annotated bibliography of the years 2011 to 2016. In *Machine Learning for Dynamic Software Analysis: Potentials and Limits, International Dagstuhl Seminar 16172*, volume 11026 of *LNCS*, 123–148. Springer.
- Jung, J. C.; Lutz, C.; Pulcini, H.; and Wolter, F. 2020. Logical separability of incomplete data under ontologies. In *Proc. of KR 2020*, 517–528.
- Kolaitis, P. G. 2011. Schema Mappings and Data Examples: Deriving Syntax from Semantics (Invited Talk). In *Proc. of FSTTCS 2011*, volume 13 of *Leibniz International Proceedings in Informatics (LIPIcs)*, 25–25. Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- Konev, B.; Lutz, C.; Ozaki, A.; and Wolter, F. 2017. Exact learning of lightweight description logic ontologies. *J. Mach. Learn. Res.* 18:201:1–201:63.
- Konev, B.; Ozaki, A.; and Wolter, F. 2016. A model for learning description logic ontologies based on exact learning. In *Proc. of AAAI*, 1008–1015. AAAI Press.
- Lehmann, J., and Hitzler, P. 2010. Concept learning in description logics using refinement operators. *Machine Learning* 78:203–250.
- Lemieux, C.; Park, D.; and Beschastnikh, I. 2015. General LTL specification mining (t). In *Proc. of ASE*, 81–92. IEEE.
- Lutz, C.; Wolter, F.; and Zakharyashev, M. 2008. Temporal description logics: A survey. In *Proc. of TIME 2008*, 3–14. IEEE Computer Society.
- Martins, D. M. L. 2019. Reverse engineering database queries from examples: State-of-the-art, challenges, and research opportunities. *Inf. Syst.* 83:89–100.

- Neider, D., and Gavran, I. 2018. Learning linear temporal properties. In *Proc. of FMCAD 2018*, 1–10. IEEE.
- Nesetril, J., and Tardif, C. 2005. Short answers to exponentially long questions: Extremal aspects of homomorphism duality. *SIAM J. Discret. Math.* 19(4):914–920.
- Schild, K. 1993. Combining terminological logics with tense logic. In *Proc. of EPIA'93*, volume 727 of *LNCS*, 105–120. Springer.
- Shahbaz, M., and Groz, R. 2009. Inferring mealy machines. In *International Symposium on Formal Methods*, 207–222. Springer.
- ten Cate, B., and Dalmau, V. 2021. Conjunctive queries: Unique characterizations and exact learnability. In *Proc. of ICDT 2021*, volume 186 of *LIPICs*, 9:1–9:24. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- ten Cate, B.; Kolaitis, P. G.; Qian, K.; and Tan, W. 2018. Active learning of GAV schema mappings. In *Proc. of PODS 2018*, 355–368. ACM.
- ten Cate, B.; Dalmau, V.; and Kolaitis, P. G. 2013. Learning schema mappings. *ACM Trans. Database Syst.* 38(4):28:1–28:31.
- Tirtarasa, S., and Turhan, A.-Y. 2022. Computing Generalizations of Temporal \mathcal{EL} Concepts with Next and Global. In *Proc. of SAC'22*. Association for Computing Machinery.