# Chain Length and CSPs Learnable with Few Queries 

Christian Bessiere, ${ }^{1}$ Clément Carbonnel, ${ }^{1}$ George Katsirelos ${ }^{2}$<br>${ }^{1}$ LIRMM, CNRS, University of Montpellier, France<br>\{christian.bessiere, clement.carbonnel\} @lirmm.fr<br>${ }^{2}$ UMR MIA-Paris, INRA, AgroParisTech, Université Paris-Saclay, Paris, France<br>gkatsi@gmail.com


#### Abstract

The goal of constraint acquisition is to learn exactly a constraint network given access to an oracle that answers truthfully certain types of queries. In this paper we focus on partial membership queries and initiate a systematic investigation of the learning complexity of constraint languages. First, we use the notion of chain length to show that a wide class of languages can be learned with as few as $O(n \log (n))$ queries. Then, we combine this result with generic lower bounds to derive a dichotomy in the learning complexity of binary languages. Finally, we identify a class of ternary languages that eludes our framework and hints at new research directions.


## 1 Introduction

Constraint programming is a declarative programming paradigm for solving combinatorial problems, with a wide range of applications in artificial intelligence (van Beek and Chen 1999; Négrevergne and Guns 2015) and operational research (Shaw 1998; Hooker and van Hoeve 2018). The user declares a list of variables, which range over a domain of his choice, and then constraints, which specify forbidden assignments to certain subsets of variables. A solver will then either find an assignment to the variables that satisfies every constraint or prove that no such assignment exists.

From the user's perspective, the difficulty lies in the modelling phase. Variables and domain values are often easily derived from the unkowns in the original problem, but specifying rigorously what a solution is using only relations imposed upon subsets of variables may require considerable expertise, especially if the solver only offers a limited cata$\log$ of possible relations. Constraint acquisition systems are learning agents designed to assist the user in modelling his problem. Such systems provide the user with simple queries - for instance membership queries, in which a given assignment must be labelled as either "solution" or "non-solution" - and use the collected data to construct a constraint network that hopefully corresponds to user's problem. There exist a number of constraint acquisition systems (Bessiere et al. 2005; 2013; Beldiceanu and Simonis 2012), with various learning strategies and theoretical guarantees.

[^0]For constraint acquisition systems, there is a tradeoff between the computational difficulty of answering the queries and the number of queries needed to identify the target network. Some queries that could guarantee quick convergence towards a constraint network are too difficult for a human to answer with certainty. For instance, equivalence queries provide the user with a candidate constraint network $N$ and ask if it captures his problem; if it does not, the user must reply with an assignment misclassified by $N$. Other queries (such as membership queries) are easy to answer but convergence may require exponentially many queries, even on very simple constraint networks (Bessiere et al. 2017).

Partial membership queries are generalisations of membership queries in which the user is required to classify partial assignments. In 2013, Bessiere et al. presented a constraint acquisition algorithm Quacq based on partial membership queries, and they observed that on certain elementary constraint languages $(\{=, \neq\}$ and $\{>\}$ on the Boolean domain) their algorithm is guaranteed to converge after $O(n \log (n))$ queries (Bessiere et al. 2013). These languages are learnable with few, simple queries, but they are not very useful for modelling. Are there more expressive languages with similar learning properties?

Our first contribution is a proof that Quacq-like algorithms can learn a considerably larger class of languages with a quasi-linear number of partial membership queries. Our argument revolves around the notion of chain length. When one constraint is removed from a constraint network, the solution set may either remain unchanged (if the constraint is redundant) or increase in size. The chain length of a constraint language corresponds roughly to the maximum number of times the size of the solution set may increase when the constraints of an $n$-variables network are removed one by one. We find that Quacq-like algorithms always learn constraint networks with $O(\mathrm{CL}(n) \cdot \log (n))$ queries, where $\mathrm{CL}(n)$ is the chain length of the language.

As an example, consider a system of homogeneous linear equations over a finite field. The solution set is a vector space. When an equation (a constraint) is removed, either the solution set remains the same, or new linearly independent solutions appear and the dimension increases. Since the dimension of the solution set is bounded by the num-
ber $n$ of variables, the latter case may only occur at most $n$ times if all the constraints are removed one by one: sets of homogeneous linear equations have linear chain length. As a consequence, they can be learned with a quasi-linear number of partial membership queries. The most general property known to imply linear chain length is the presence of a (finite-domain) Mal'tsev embedding, which was introduced recently in the context of parameterised complexity (Lagerkvist and Wahlström 2017). These languages can be fairly expressive; in fact, some NP-hard languages have a Mal'tsev embedding (Lagerkvist and Wahlström 2017).

As our second contribution, we show that every binary constraint language that does not have linear chain length (via a notion closely related to Mal'tsev embeddings, which we call Mal'tsev derivations) has VC-dimension $\Omega\left(n^{2}\right)$. This implies that a binary language is either learnable with $O(n \log (n))$ partial membership queries, or it is not learnable with $o\left(n^{2}\right)$ yes/no queries of any kind.

Third and last, we show that there exist ternary constraint languages with VC-dimension $o\left(n^{2}\right)$ but no Mal'tsev embedding nor Mal'tsev derivation. This suggests that classifying the learning complexity of ternary languages requires tools beyond VC-dimension and Mal'tsev derivations.

## 2 Preliminaries

CSP. A (constraint) language is a finite set of relations over a finite domain $D$. Given a language $\Gamma$, a CSP instance over $\Gamma$ is a pair $(X, C)$ where $X$ is a set of $n$ variables and $C$ is a set of constraints, that is, pairs $c=\left(R_{c}, S_{c}\right)$ where $R_{c} \in \Gamma$ and $S_{c}$ is a tuple of variables whose length is the arity of $R_{c}$. An assignment $\phi: X \rightarrow D$ satisfies a constraint $c=\left(R_{c}, S_{c}\right)$ if $\phi\left(S_{c}\right) \in R_{c}$; otherwise it violates $c$. A solution to a CSP instance is an assignment $\phi: X \rightarrow D$ that satisfies every constraint. We will assume that the variable set $X$ has a total order $x_{1}<x_{2}<\ldots<x_{n}$ and define $\operatorname{sol}(I)$ as the $n$-ary relation $\left\{\left(\phi\left(x_{1}\right), \ldots, \phi\left(x_{n}\right)\right) \mid\right.$ $\phi$ is a solution to $I\}$ over $D$. The restriction of a CSP instance $I=(X, C)$ to a subset $X^{\prime}$ of $X$, denoted by $I\left[X^{\prime}\right]$, is obtained by dropping all variables in $X \backslash X^{\prime}$ and all constraints $c$ such that $S_{c}$ contains a variable in $X \backslash X^{\prime}$.

Constraint acquisition. We use Angluin's exact learning model (Angluin 1988). Let $\Gamma$ be a language over a domain $D$. We associate to each CSP instance $I=(X, C)$ over $\Gamma$ an oracle which answers truthfully a certain type of queries. A membership query asks whether an assignment $\phi: X \rightarrow D$ is a solution to $I$, a partial membership query asks whether a partial assignment $\phi: X^{\prime} \rightarrow D$ is a solution to $I\left[X^{\prime}\right]$, and an equivalence query asks whether a CSP instance $I^{\prime}$ satisfies $\operatorname{sol}(I) \neq \operatorname{sol}\left(I^{\prime}\right)$, in which case the oracle must return a tuple in the symmetric difference of $\operatorname{sol}(I)$ and $\operatorname{sol}\left(I^{\prime}\right)$. A query is Boolean if the oracle answer can only be yes or no. Membership queries and partial membership queries are Boolean, but not equivalence queries.

A constraint acquisition algorithm for $\Gamma$ is an algorithm that, given a variable set $X$ and access to an oracle for a target CSP instance $I_{T}=\left(X, C_{T}\right)$ over $\Gamma$, outputs a CSP instance $I_{L}$ over $\Gamma$ such that $\operatorname{sol}\left(I_{T}\right)=\operatorname{sol}\left(I_{L}\right)$. Given a
type $Q$ of queries, we say that a finite constraint language $\Gamma$ is learnable with $f(n)$ queries of type $Q$ if there exists a constraint acquisition algorithm that learns any CSP instance over $\Gamma$ of $n$ variables with at most $f(n)$ queries of type $Q$. Note that the time complexity of the acquisition algorithm needs not be bounded by a polynomial.
c-definitions. Given a language $\Gamma$ over a domain $D$ and a relation $R \subseteq D^{r}$, we say that $R$ is $c$-definable over $\Gamma$ if there exists a CSP instance $I$ over $\Gamma$ such that $\operatorname{sol}(I)=R$. (c-definitions are particular cases of the more widely used qfpp-definitions, which allow equality relations in addition to the relations of $\Gamma$; see e.g. (Jonsson, Lagerkvist, and Roy 2017).) By extension, a language $\Gamma^{\prime}$ is c-definable in $\Gamma$ if every $R \in \Gamma^{\prime}$ is c-definable in $\Gamma$. We define the chain length of $\Gamma$ as the function $\mathrm{CL}_{\Gamma}: \mathbb{N}^{*} \rightarrow \mathbb{N}^{*}$ such that $\mathrm{CL}_{\Gamma}(n)$ is the maximum length of a sequence $R_{1}, R_{2}, \ldots, R_{q}$ of relations of the same arity $k \leq n$ such that

$$
R_{1} \subsetneq R_{2} \subsetneq \ldots \subsetneq R_{q}
$$

and each $R_{i}, i \in\{1, \ldots, q\}$, is c-definable over $\Gamma$.
We say that $\Gamma$ shatters a relation $R$ if for every subrelation $R^{\prime} \subseteq R$ there exists a relation $R^{*}$ such that $R^{*} \cap R=R^{\prime}$ and $R^{*}$ is c-definable over $\Gamma$. (Note that $R$ itself needs not be c-definable.) The $V C$-dimension of $\Gamma$ is the function $\mathrm{VC}_{\Gamma}: \mathbb{N}^{*} \rightarrow \mathbb{N}^{*}$ such that $\mathrm{VC}_{\Gamma}(n)$ is the maximum size (i.e. number of tuples) of a relation of arity at most $n$ that is shattered by $\Gamma$. We will also discuss a measure of redundancy, which is tightly connected to the VC-dimension of a language. Given a CSP instance $(X, C)$, we say that a constraint $c \in C$ is non-redundant if there exists an assignment $\phi: X \rightarrow D$ such that $\phi$ is a solution to $(X, C \backslash\{c\})$ but $\phi$ violates $c$. The non-redundancy of a language $\Gamma$, which we denote by $\operatorname{NRD}_{\Gamma}(n)$, is the maximum number of nonredundant constraints in a CSP instance over $\Gamma$ with $n$ variables. Throughout the paper we will use the following observations about $\mathrm{CL}_{\Gamma}(n), \mathrm{VC}_{\Gamma}(n)$ and $\mathrm{NRD}_{\Gamma}(n)$.
Observation 1. If $\Gamma=\Gamma_{1} \cup \Gamma_{2}$, then $C L_{\Gamma}(n) \leq C L_{\Gamma_{1}}(n)+$ $C L_{\Gamma_{2}}(n)$.

Proof. Let $n>0, q=\mathrm{CL}_{\Gamma}(n)$ and $R_{1}, \ldots, R_{q}$ be a sequence of relations of arity at most $n$, c-definable over $\Gamma$, and such that $R_{1} \subsetneq R_{2} \subsetneq \ldots \subsetneq R_{q}$. Let $I_{1}, \ldots, I_{q}$ be CSP instances over $\Gamma$ such that $I_{j}=\left(X, C_{j}\right)$ and $\operatorname{sol}\left(I_{j}\right)=R_{j}$ for all $j \leq q$. We assume without loss of generality that $C_{j+1} \subsetneq C_{j}$ for all $j<q$. (If it is not the case, we add to $C_{j}$ all constraints in $\cup_{k>j} C_{k}$; this operation leaves $\operatorname{sol}\left(I_{j}\right)$ unchanged.) We partition each constraint set $C_{j}$ into $C_{j}^{1}=$ $\left\{(R, S) \in C_{j} \mid R \in \Gamma_{1}\right\}, C_{j}^{2}=\left\{(R, S) \in C_{j} \mid R \in \Gamma_{2}\right\}$ and define $R_{j}^{1}=\operatorname{sol}\left(X, C_{j}^{1}\right), R_{j}^{2}=\operatorname{sol}\left(X, C_{j}^{2}\right)$. Then, we have $R_{j}=R_{j}^{1} \cap R_{j}^{2}$ for all $j \leq q$. Furthermore, since $R_{1} \subsetneq \ldots \subsetneq R_{q}, R_{1}^{1} \subseteq \ldots \subseteq R_{q}^{1}$ and $R_{1}^{2} \subseteq \ldots \subseteq R_{q}^{2}$, it follows that for all $j<q$ we have either $R_{j}^{1} \subsetneq R_{j+1}^{1}$ or $R_{j}^{2} \subsetneq R_{j+1}^{2}$. These events may happen at most $\mathrm{CL}_{\Gamma_{1}}(n)$ times and $\mathrm{CL}_{\Gamma_{2}}(n)$ times, respectively. This implies $q=$ $\mathrm{CL}_{\Gamma}(n) \leq \mathrm{CL}_{\Gamma_{1}}(n)+\mathrm{CL}_{\Gamma_{2}}(n)$.

Observation 2. If $\Gamma$ is $c$-definable over $\Gamma^{\prime}$, then $C L_{\Gamma}(n) \leq$ $C L_{\Gamma^{\prime}}(n)$.

Proof. The claim follows directly from the fact that every relation c-definable over $\Gamma$ is also c-definable over $\Gamma^{\prime}$.

Observation 3. If $\Gamma$ has VC-dimension $\Omega(f(n))$, then $\Gamma$ is not learnable with o $(f(n))$ Boolean queries.

Proof. Suppose that $A$ is a constraint acquisition algorithm that learns any CSP over $\Gamma$ with $o(f(n))$ queries. The output of $A$ is completely determined by the $o(f(n))$ bits that it receives as answers to its Boolean queries. There are $2^{o(f(n))}$ possible combinations of inputs, so the algorithm can only output $2^{o(f(n))}$ different CSP instances over $n$ variables. However, since $\Gamma$ has VC-dimension $\Omega(f(n))$, the number of CSP instances over $\Gamma$ with distinct solution sets and $n$ variables is $2^{\Omega(f(n)}$. It follows that for sufficiently large $n$ the algorithm has fewer outputs than there are distinct $n$ variables constraint networks over $\Gamma$, a contradiction.

Observation 4. For every language $\Gamma$, it holds that $N R D_{\Gamma}(n)=V C_{\Gamma}(n)$.

Proof. We first prove $\mathrm{NRD}_{\Gamma}(n) \leq \mathrm{VC}_{\Gamma}(n)$. Let $n>0$ and $I=(X, C)$ be a CSP instance over $\Gamma$ with $n$ variables and such that $C$ contains precisely $\operatorname{NRD}_{\Gamma}(n)$ constraints, all of which are non-redundant. Then, for each $c \in C$ there exists an assignment $\phi_{c}$ such that $\phi_{c}$ violates $c$ but satisfies every other constraint. For each $c \in C$, let $t_{c}$ be a tuple such that $t_{c}[i]=\phi_{c}\left(x_{i}\right)$ for all $i \in\{1, \ldots, n\}$. We claim that $\Gamma$ shatters the relation $R_{I}=\left\{t_{c} \mid c \in C\right\}$. For each $R_{I}^{*} \subseteq R_{I}$, we define the instance $I^{*}=\left(X, C^{*}\right)$ with $\left.\left\{c \in C \mid t_{c} \notin R_{I}^{*}\right\}\right)$. Then, $\phi_{c}$ is a solution to $I^{*}$ if and only if $c \notin C^{*}$, and hence $R_{I} \cap \operatorname{sol}\left(I^{*}\right)=R_{I}^{*}$. Therefore, $\Gamma$ shatters $R_{I}$, which contains $\operatorname{NRD}_{\Gamma}(n)$ tuples, and hence $\mathrm{NRD}_{\Gamma}(n) \leq \mathrm{VC}_{\Gamma}(n)$. Now we prove $\operatorname{NRD}_{\Gamma}(n) \geq \mathrm{VC}_{\Gamma}(n)$. Let $R$ be a relation of arity $r \leq n$ such that $\Gamma$ shatters $R$ and $q=|R|=\mathrm{VC}_{\Gamma}(n)$. Let $t_{1}, \ldots, t_{q}$ be an arbitrary ordering of the tuples of $R$ and $I_{1}, \ldots, I_{q}$ be a set of CSP instances over $\Gamma$ with variable set $X$ and constraint sets $C_{i}, 1 \leq i \leq q$, such that $\operatorname{sol}\left(I_{i}\right) \cap R=R \backslash\left\{t_{i}\right\}$. Note that each $C_{i}$ can be chosen so that it contains a single constraint $c_{i}$, since $t_{i} \notin \operatorname{sol}\left(I_{i}\right)$ if and only if there exists $c_{i} \in C_{i}$ such that $t_{i} \notin \operatorname{sol}\left(X,\left\{c_{i}\right\}\right)$. This implies that no constraint $c$ may belong to $C_{i}$ and $C_{j}$ for any $j \neq i$, since otherwise we would have $\operatorname{sol}\left(I_{i}\right)=\operatorname{sol}\left(I_{j}\right)$. If we define $I^{*}=\left(X, C_{1} \cup \ldots \cup C_{q}\right)$ then this instance has exactly $q=\mathrm{VC}_{\Gamma}(n)$ constraints $c_{1}, \ldots, c_{q}$. Furthermore, removing any constraint $c_{i}$ adds at least the tuple $t_{i}$ to $\operatorname{sol}\left(I^{*}\right)$, so no constraint is redundant and $\operatorname{NRD}_{\Gamma}(n) \geq \mathrm{VC}_{\Gamma}(n)$.

Observation 5. For every language $\Gamma$, it holds that $V C_{\Gamma}(n) \leq \sum_{R \in \Gamma} V C_{\{R\}}(n)$.
Proof. By Observation 4 it suffices to show that every CSP instance $I=(X, C)$ over $\Gamma$ with $n$ variables has at most $\sum_{R \in \Gamma} \operatorname{NRD}_{\{R\}}(n)$ non-redundant constraints. For each $R \in \Gamma$ we define $I_{R}=\left(X,\left\{c=\left(R, S_{c}\right) \in C\right\}\right)$. Note that if $c=(R, S)$ is non-redundant in $I$, then it is nonredundant in $I_{R}$. Each instance $I_{R}$ has at most $\operatorname{NRD}_{\{R\}}(n)$ non-redundant constraints, so the claim follows.

Partial polymorphisms. A partial operation of arity $r$ on a set $D$ is a mapping $f: Q \rightarrow D$ where $Q \subseteq D^{r}$. The set $Q$ is called the domain of $f$ and we write $Q=\operatorname{domain}(f)$. Given a relation $R$ of arity $k$ over a set $D$ and a partial operation $f$ on $D$ of arity $r$, we say that $f$ is a partial polymorphism of $R$ if for every $t_{1}, \ldots, t_{r} \in R$, either $f\left(t_{1}, \ldots, t_{r}\right) \in R$ or there exists $i \in\{1, \ldots, k\}$ such that $\left(t_{1}[i], \ldots, t_{r}[i]\right) \notin \operatorname{domain}(f)$. If in addition we have domain $(f)=D^{r}$, we say that $f$ is a polymorphism of $R$. A partial operation is a (partial) polymorphism of a language $\Gamma$ if it is a (partial) polymorphism of each of its relations. Polymorphisms and partial polymorphisms are fundamental technical tools in the study of the complexity and expressivity of constraint languages. For a comprehensive exposition of the subject we refer the reader to (Barto, Krokhin, and Willard 2017; Couceiro, Haddad, and Lagerkvist 2019).

## 3 Constraint acquisition and chain length

In this section we describe a simple constraint acquisition algorithm that learns any CSP instance with $n$ variables over a language $\Gamma$ using $O\left(\mathrm{CL}_{\Gamma}(n) \cdot \log (n)\right)$ partial membership queries. Our contribution does not lie in the algorithm itself, which is essentially a stripped-down version of the constraint acquisition algorithm Quacq found in (Bessiere et al. 2013), but rather in its analysis. We borrow the following proposition from their work.
Proposition 6 ((Bessiere et al. 2013), Proposition 2). There exists an algorithm FindScope which given in input

- a set $X$ of $n$ variables,
- a partial membership oracle for a target CSP instance $\left(X, C_{T}\right)$, and
- a negatively answered membership query $\phi$
outputs the set $X_{c}$ of variables that appear in the scope $S_{c}$ of a constraint $c \in C_{T}$ that $\phi$ violates, using $O\left(\left|S_{c}\right| \cdot \log (n)\right)$ partial membership queries.
Theorem 7. Let $\Gamma$ be a fixed language. There exists a constraint acquisition algorithm that learns any CSP instance over $\Gamma$ on $n$ variables using $O\left(C L_{\Gamma}(n) \cdot \log (n)\right)$ partial membership queries.

Proof. Let $O$ be an oracle associated with a CSP instance $I_{T}=\left(X, C_{T}\right)$ over $\Gamma$. Consider the following constraint acquisition algorithm. Let $L$ and $P$ be two sets of constraints from the language $\Gamma$ on the variable set $X$. These sets can be understood as the sets of learned and possible constraints, respectively. When the algorithm is invoked, $L$ is initialised to the empty set and $P$ is initialised to the set of all possible constraints from $\Gamma$ on $X$. The algorithm then repeats the following steps until $\operatorname{sol}(X, L) \subseteq \operatorname{sol}(X, P)$ :

- Find $\phi \in \operatorname{sol}(X, L) \backslash \operatorname{sol}(X, P)$.
- Query $\phi$ to the oracle.
- If the answer is yes, set $P=\{c \in P: \phi$ satisfies $c\}$.
- If the answer is no, invoke FindScope on input $(X, O, \phi)$, and then query all possible partial assignments to the returned subset $X^{\prime}$ to determine $\operatorname{sol}\left(I_{T}\left[X^{\prime}\right]\right)$. Add to $L$ a set $C^{\prime}$ of constraints such that $\operatorname{sol}\left(X^{\prime}, C^{\prime}\right)=\operatorname{sol}\left(I_{T}\left[X^{\prime}\right]\right)$.

Once the condition $\operatorname{sol}(X, L) \subseteq \operatorname{sol}(X, P)$ holds, the algorithm returns $(X, L)$ and halts.

We now turn to the analysis. First, observe that at each loop either $L$ grows or $P$ shrinks, so the algorithm will eventually halt. For correctness, it suffices to show that we have the two invariants (i) $\operatorname{sol}\left(X, C_{T}\right) \subseteq \operatorname{sol}(X, L)$ and (ii) $C_{T} \subseteq P$. (If both (i) and (ii) hold, then $\operatorname{sol}(X, P) \subseteq$ $\operatorname{sol}\left(X, C_{T}\right) \subseteq \operatorname{sol}(X, L)$ is also an invariant, and upon termination we have $\operatorname{sol}\left(X, C_{T}\right)=\operatorname{sol}(X, L)$.) Both (i) and (ii) hold at initialisation. Notice that whenever a constraint $c$ is removed from $P$, there exists a solution $\phi$ to $\left(X, C_{T}\right)$ that violates $c$, so $c \notin C_{T}$ and (ii) is indeed an invariant of the algorithm. Similarly, whenever a constraint $c$ is added to $L$ we have $\operatorname{sol}\left(X, C_{T} \cup\{c\}\right)=\operatorname{sol}\left(X, C_{T}\right)$; it follows that if $\operatorname{sol}\left(X, C_{T}\right) \subseteq \operatorname{sol}(X, L)$ then $\operatorname{sol}\left(X, C_{T}\right)=$ $\operatorname{sol}\left(X, C_{T} \cup\{c\}\right) \subseteq \operatorname{sol}(X, L \cup\{c\})$, so (i) is also an invariant and the algorithm is correct.

Now, we bound the number of oracle queries performed by the algorithm before it halts. We will use $L_{i}, 0 \leq i \leq l$, and $P_{j}, 0 \leq j \leq p$ to denote the sequences of distinct states of $L$ and $P$, respectively. Using Proposition 6 to bound the number of oracle queries within each call to FindScope, and observing that any subset $X^{\prime}$ of variables returned by FindScope has bounded size since $\Gamma$ (and hence the maximum constraint arity) is fixed, we find that the algorithm makes a total of $O(p+l \cdot \log (n))$ oracle queries. Thus, it suffices to bound $l$ and $p$ by a linear function of $\mathrm{CL}_{\Gamma}(n)$.

First, observe that for every $i \in\{0, \ldots, l-1\}$, it holds that $\operatorname{sol}\left(X, L_{i+1}\right) \subsetneq \operatorname{sol}\left(X, L_{i}\right)$ since $L_{i} \subsetneq L_{i+1}$ and $\operatorname{sol}\left(X, L_{i+1}\right)\left[X^{\prime}\right]=\operatorname{sol}\left(X, C_{T}\right)\left[X^{\prime}\right] \neq \operatorname{sol}\left(X, L_{i}\right)\left[X^{\prime}\right]$, where $X^{\prime}$ is the set of variables returned by FindScope. Each relation $\operatorname{sol}\left(X, L_{i}\right), 0 \leq i \leq l$, is c-definable over $\Gamma$ so we deduce $l \leq \mathrm{CL}_{\Gamma}(n)$.

In order to derive a similar bound on $p$, let $\phi_{j}, 0 \leq j \leq$ $p-1$, be the positively answered query that triggered the update of $P$ from state $P_{j}$ to state $P_{j+1}$, and let $c_{j}$ be an arbitrary constraint in $P_{j} \backslash P_{j+1}$. In particular, observe that the assignment $\phi_{j}$ violates $c_{j}$, but $\phi_{j}$ does not violate $c_{k}$ for any $k>j$. Thus, we have

$$
\begin{aligned}
& \operatorname{sol}\left(X,\left\{c_{p}\right\}\right) \subsetneq \operatorname{sol}\left(X,\left\{c_{p-1}, c_{p}\right\}\right) \\
& \subsetneq \ldots \subsetneq \operatorname{sol}\left(X,\left\{c_{0}, \ldots, c_{p}\right\}\right)
\end{aligned}
$$

since for each $0 \leq j \leq p-1$ the assignment $\phi_{j}$ is a solution to $\left(X,\left\{c_{j+1}, \ldots, c_{p}\right\}\right)$ but not to ( $X,\left\{c_{j}, \ldots, c_{p}\right\}$ ). Again, each of these relations is c-definable over $\Gamma$ so $p \leq \mathrm{CL}_{\Gamma}(n)$. The total query complexity of the algorithm is $O\left(\mathrm{CL}_{\Gamma}(n)+\right.$ $\left.\mathrm{CL}_{\Gamma}(n) \cdot \log (n)\right)=O\left(\mathrm{CL}_{\Gamma}(n) \cdot \log (n)\right)$, as claimed.

## 4 Languages of linear chain length

As a corollary to Theorem 7, every language $\Gamma$ with linear chain length is learnable with a quasi-linear number of partial membership queries. As noted in the introduction, linear equations over finite fields have linear chain length. They correspond to a special case of a larger class of languages; to describe this class we will need additional terminology.

A Mal'tsev operation on a domain $D$ is a ternary operation $f: D^{3} \rightarrow D$ such that $f(x, x, y)=f(y, x, x)=y$ for all $x, y \in D$. A language $\Gamma$ is Mal'tsev if it has a Mal'tsev
polymorphism. Mal'tsev languages have played an important role in the study of the complexity of constraint languages (Bulatov and Dalmau 2006; Idziak et al. 2010). Their most salient property is that any c-definable relation has a linear-size generating set which encodes them efficiently, in the same way that bases describe vector spaces. The fact that Mal'tsev languages have linear chain length is folklore; we include a proof for completeness.

The signature of a relation $R$ over $D$ of arity $r$, denoted by $\operatorname{Sig}(R)$, is the set of all triples $\left(k, d_{1}, d_{2}\right) \in\{1, \ldots, r\} \times D^{2}$ such that there exist $t_{1}, t_{2} \in R$ satisfying $t_{1}[k]=d_{1}, t_{2}[k]=$ $d_{2}$ and $t_{1}[i]=t_{2}[i]$ for all $i<k$.
Lemma 8 ((Bulatov and Dalmau 2006), Lemma 3.1). Let $\phi$ be a Mal'tsev operation over a domain $D$ and $R_{1}, R_{2}$ be two relations over $D$ such that $R_{1} \subseteq R_{2}$ and $\phi$ is a polymorphism of both $R_{1}$ and $R_{2}$. Then, $\operatorname{Sig}\left(R_{1}\right)=\operatorname{Sig}\left(R_{2}\right)$ if and only if $R_{1}=R_{2}$.

Proposition 9. If $\Gamma$ is Mal'tsev, then $C L_{\Gamma}(n)=O(n)$.

Proof. Let $\phi$ be a Mal'tsev polymorphism of $\Gamma$. Let $R_{1}, \ldots, R_{q}$ be a sequence of $n$-ary relations, $c$-definable over $\Gamma$, such that $R_{1} \subsetneq \ldots \subsetneq R_{q}$. Polymorphisms are invariant under $c$-definitions (Barto, Krokhin, and Willard 2017), so $\phi$ is a polymorphism of $R_{1}, \ldots, R_{q}$. By Lemma 8 we have $\operatorname{Sig}\left(R_{1}\right) \subsetneq \ldots \subsetneq \operatorname{Sig}\left(R_{q}\right)$. Since $\left|\operatorname{Sig}\left(R_{q}\right)\right| \leq n|D|^{2}$, we obtain $q \leq n|D|^{2}+1=O(n)$ and $\mathrm{CL}_{\Gamma}(n)=O(n)$.

We can extend this result using Observation 1 and Observation 2, which provide ways to combine languages and relations while preserving chain length. We say that a constraint language $\Gamma$ has a Mal'tsev derivation if there exists a language $\hat{\Gamma}$ such that (i) every relation in $\hat{\Gamma}$ has a Mal'tsev polymorphism and (ii) every relation in $\Gamma$ is c-definable in $\hat{\Gamma}$. (Intuitively, languages with Mal'tsev derivations are those that can be obtained from a finite number of Mal'tsev relations through a finite sequence of language unions and cdefinitions.) The following proposition follows from Proposition 9, Observation 1 and Observation 2.

Proposition 10. Let $\Gamma$ be a language. If $\Gamma$ has a Mal'tsev derivation, then $C L_{\Gamma}(n)=O(n)$.

Combining this result with Theorem 7, we obtain a general sufficient condition for learnability with a quasi-linear number of partial membership queries.

Theorem 11. If $\Gamma$ has a Mal'tsev derivation, then $\Gamma$ is learnable with $O(n \log (n))$ partial membership queries.

Example 1. A classical example of Mal'tsev relations are those corresponding to the solution space of an equation over a finite field (see e.g. Example 5 in (Bulatov and Dalmau 2006)). Theorem 11 implies that any language in which each individual relation is the solution space of an equation over a finite field is learnable with a quasi-linear number of partial membership queries, even if the field in question differs from one relation to another.

Example 2. Let $k \geq 3$ be a fixed integer and consider the relation $R_{k}$ that contains all Boolean $k$-tuples of Hamming weight 1 . This relation is not Mal'tsev as mappings satisfying the Mal'tsev identity must map any three tuples of $R_{k}$ to one with Hamming weight at least 2 . However, $R_{k}$ is precisely the intersection of $\{0,1\}^{k}$ with the solution set $R^{*}$ of the equation $\left(x_{1}+\ldots+x_{k}\right) \bmod k=1$. Both $R^{*}$ and $\{0,1\}^{k}$ are Mal'tsev (we invite the reader to verify that the operation $f(x, y, z)=x-y+z \bmod k$ is a polymorphism of $R^{*}$ ), so $\left\{R_{k}\right\}$ has a Mal'tsev derivation.

Mal'tsev derivations are very similar to a notion introduced in (Lagerkvist and Wahlström 2017) called Mal'tsev embeddings. A language $\Gamma$ over $D$ embeds in a language $\hat{\Gamma}$ over $\hat{D} \supset D$ if there exists a bijective function $h: \Gamma \rightarrow \hat{\Gamma}$ such that for each $R \in \Gamma$ of arity $r$, we have that $R=$ $h(R) \cap D^{r}$. If $\Gamma$ embeds in a Mal'tsev language whose domain is finite, we say that $\Gamma$ has a (finite) Mal'tsev embedding. For instance, in Example 2 the language $\left\{R_{k}\right\}$ embeds in $\left\{R^{*}\right\}$, which is Mal'tsev.

Every language with a Mal'tsev embedding has a straightforward Mal'tsev derivation: if $\Gamma$ is a language over a domain $D$ that embeds into a Mal'tsev language $\Gamma^{\prime}$ of arity at most $r$, then it has a c-definition over $\hat{\Gamma}=\Gamma^{\prime} \cup\left\{D^{k}: k \leq r\right\}$, in which each relation is Mal'tsev. Could the two notions coincide? Building on ideas from (Lagerkvist and Wahlström 2017), we can deduce a partial converse.

Observation 12. If $\Gamma$ has a Mal'tsev derivation, then it has an infinite-domain Mal'tsev embedding.
Proof sketch. Theorem 28 in (Lagerkvist and Wahlström 2017) can be generalised to arbitrary domains to show that (i) infinite-domain Mal'tsev embeddings can be characterised in terms of partial polymorphisms only, and (ii) for any two languages $\Gamma_{1}, \Gamma_{2}$ with Mal'tsev embeddings, there exists an infinite domain $D_{\infty}$ and a Mal'tsev operation $u$ over $D_{\infty}$ such that both $\Gamma_{1}$ and $\Gamma_{2}$ can be embedded in a language over $D_{\infty}$ with polymorphism $u$. From (i) and (ii), we obtain that infinite-domain Mal'tsev embeddings are invariant under c-definitions and language unions.

## 5 A dichotomy for binary languages

The results of Sections 3 and 4 raise interesting questions. Are there languages learnable with a quasi-linear number of partial membership queries but whose chain length is not linear? Is there a gap, in the sense that a language not learnable with $O(n \log (n))$ Boolean queries is not learnable with $o\left(n^{c}\right)$ Boolean queries of any kind, for some $c>1$ ?

In this section we prove the following theorem, which answers both questions for binary languages. A binary relation $R$ is rectangular if there does not exist four (not necessarily distinct) values $a, b, c, d$ such that $(a, b),(c, d),(a, d) \in R$ but $(c, b) \notin R$. A language is rectangular if each of its relations is rectangular.
Theorem 13. Let $\Gamma$ be a binary language. If $\Gamma$ is rectangular, then it has linear chain length and is learnable with $O(n \log (n))$ partial membership queries. Otherwise, $\Gamma$ is not learnable with o( $n^{2}$ ) Boolean queries.

We first prove that every rectangular relation has a Mal'tsev derivation. To this purpose, we will need two known facts about rectangular relations.
Lemma 14 ((Kazda 2011), Lemma 10). Let $R$ be a rectangular relation. If we define the relations $\sim_{1}, \sim_{2}$ over $D$ as

$$
\begin{aligned}
& d \sim_{1} d^{\prime} \Longleftrightarrow \exists d^{*} \in D \mid\left(d, d^{*}\right) \in R \text { and }\left(d^{\prime}, d^{*}\right) \in R \\
& d \sim_{2} d^{\prime} \Longleftrightarrow \exists d^{*} \in D \mid\left(d^{*}, d\right) \in R \text { and }\left(d^{*}, d^{\prime}\right) \in R
\end{aligned}
$$

then $\sim_{1}$ and $\sim_{2}$ are equivalence relations on the sets $\{d \in$ $\left.D \mid \exists d^{*} \in D,\left(d, d^{*}\right) \in R\right\}$ and $\left\{d \in D \mid \exists d^{*} \in\right.$ $\left.D,\left(d^{*}, d\right) \in R\right\}$, respectively.

Furthermore, if we denote by $E_{1}, E_{2}$ the equivalence classes of $\sim_{1}, \sim_{2}$ then the mapping $h: E_{1} \rightarrow E_{2}$ given by

$$
h(e)=f \Longleftrightarrow \exists d_{e} \in e, d_{f} \in f:\left(d_{e}, d_{f}\right) \in R
$$

is a bijection from $E_{1}$ to $E_{2}$.
Lemma 15 ((Kazda 2011), Lemma 15). Let $R$ be a rectangular relation and $\sim_{1}$ be as in Lemma 14. Let $R^{+}$be the relation whose domain is the set $E_{1}$ of equivalence classes of $\sim_{1}$, and such that

$$
(e, f) \in R^{+} \Longleftrightarrow \exists d_{1} \in e, d_{2} \in f:\left(d_{1}, d_{2}\right) \in R
$$

Then, $R$ is Mal'tsev if and only if $R^{+}$is Mal'tsev.
Proposition 16. Let $\Gamma$ be a binary language. If $\Gamma$ is rectangular, then it has a Mal'tsev derivation.

Proof. Let $D$ be the domain of $\Gamma$. We will prove that the language $\{R\}$ embeds in a Mal'tsev language $\{\hat{R}\}$ for every $R \in \Gamma$. The claim will then follow from the fact that $\Gamma$ has a c-definition over $\hat{\Gamma}=\{\hat{R}: R \in \Gamma\} \cup\left\{D^{2}\right\}$, in which every relation is Mal'tsev. Let $R$ be a (rectangular) relation in $\Gamma$ and $\sim_{1}, \sim_{2}, E_{1}, E_{2}, h$ be as in Lemma 14. For each ordered pair $(e, f) \in\left(E_{1}\right)^{2}$ we introduce a fresh domain value $\alpha_{e f}$. Then, we define

$$
\begin{aligned}
& S_{1}=\left\{\left(\alpha_{e f}, d^{*}\right) \mid(e, f) \in\left(E_{1}\right)^{2}, d^{*} \in h(e)\right\} \\
& S_{2}=\left\{\left(d^{*}, \alpha_{e f}\right) \mid(e, f) \in\left(E_{1}\right)^{2}, d^{*} \in f\right\} \\
& S_{3}=\left\{\left(\alpha_{e f}, \alpha_{g e}\right) \mid(e, f, g) \in\left(E_{1}\right)^{3}\right\}
\end{aligned}
$$

and a relation $\hat{R}=R \cup S_{1} \cup S_{2} \cup S_{3}$, whose domain is denoted by $\hat{D}$. By construction, $\hat{R} \cap D^{2}=R$.

We claim that $\hat{R}$ is rectangular. It follows from the definitions of $S_{1}, S_{2}$ and $S_{3}$ that for any $\alpha_{e f} \in \hat{D} \backslash D, \hat{d} \in \hat{D}$ and $d_{e} \in e,\left(\alpha_{e f}, \hat{d}\right)$ is a tuple of $\hat{R}$ if and only if $\left(d_{e}, \hat{d}\right)$ is a tuple of $\hat{R}$. Symmetrically, for any $\alpha_{e f} \in \hat{D} \backslash D, \hat{d} \in \hat{D}$ and $d_{h(f)} \in h(f),\left(\hat{d}, \alpha_{e f}\right)$ is a tuple of $\hat{R}$ if and only if $\left(\hat{d}, d_{h(f)}\right)$ is a tuple of $\hat{R}$. Towards a contradiction, let $(a, b, c, d) \in \hat{D}$ be such that $(a, b),(c, d),(a, d) \in \hat{R}$ but $(c, b) \notin \hat{R}$. By the observation above, if any of $a, b, c, d$ belongs to $\hat{D} \backslash D$ then it can be replaced by a value in $D$ while still having $(a, b),(c, d),(a, d) \in \hat{R}$ and $(c, b) \notin \hat{R}$. If we replace them all, then we obtain values $(a, b, c, d) \in D$ such
that $(a, b),(c, d),(a, d) \in \hat{R}$ but $(c, b) \notin \hat{R}$, which is impossible since $R=\hat{R} \cap D^{2}$ is rectangular.

Since $\hat{R}$ is rectangular, we can define the relation $\hat{\sim}_{1}$ as in Lemma 14 and $\hat{R}^{+}$as in Lemma 15, whose domain is the set $\hat{E}$ of equivalence classes of $\hat{\sim}_{1}$. By construction we have

$$
\hat{E}=\left\{e \cup\left\{\alpha_{e f} \mid f \in E\right\} \mid e \in E\right\}
$$

and $\left(\alpha_{e f}, \alpha_{f e}\right) \in \hat{R}$ for all $f, e \in E$. Then, for all $\hat{e}, \hat{f} \in \hat{E}$ we have $(\hat{e}, \hat{f}) \in \hat{R}^{+}$since $\left(\alpha_{e f}, \alpha_{f e}\right) \in \hat{R}, \alpha_{e f} \in \hat{e}$ and $\alpha_{f e} \in \hat{f}$. Therefore $\hat{R}^{+}=\hat{E}^{2}$, which is Mal'tsev as it admits all operations on $\hat{E}$ as polymorphisms. By Lemma 15, $\hat{R}$ is Mal'tsev as well, and hence $\{R\}$ has a Mal'tsev embedding. This is true for all $R \in \Gamma$, so $\Gamma$ has a Mal'tsev derivation over $\hat{\Gamma}=\{\hat{R}: R \in \Gamma\} \cup\left\{D^{2}\right\}$.

Proposition 17. Let $\Gamma$ be a binary language. If $\Gamma$ is not rectangular, then $V C_{\Gamma}(n)=\Omega\left(n^{2}\right)$.

Proof. Let $R \in \Gamma$ be a non-rectangular relation. By definition, there exist three tuples $(a, b),(c, d),(a, d) \in R$ such that $(c, b) \notin R$. Now, let $n$ be an even positive integer. For each pair $(i, j)$ with $1 \leq i \leq n / 2$ and $n / 2<j \leq n$, let $t_{(i, j)}$ denote the $n$-tuple given by $t[i]=c, t[j]=b$, and for all $k \in$ $\{1, \ldots, n\}$ different from $i, j$ we have $t[k]=a$ if $k \leq n / 2$ and $t[k]=d$ otherwise. We claim that $\Gamma$ shatters the relation $R_{n}=\left\{t_{(i, j)} \mid(i, j) \in\{1, \ldots, n / 2\} \times\{n / 2+1, \ldots, n\}\right\}$.

Let $R_{n}^{*}$ be an arbitrary subrelation of $R_{n}$. Let $I^{*}=$ $\left(X, C^{*}\right)$ be a CSP instance such that $C^{*}=\left\{\left(R,\left(x_{i}, x_{j}\right)\right) \mid\right.$ $\left.t_{(i, j)} \notin R_{n}^{*}\right\}$. By construction, an assignment $\phi$ to $X$ such that $\phi\left(x_{k}\right) \in\{a, c\}$ for $k \leq n / 2$ and $\phi\left(x_{k}\right) \in\{b, d\}$ otherwise is a solution to $I^{*}$ if and only if there is no constraint $\left(R,\left(x_{i}, x_{j}\right)\right) \in C^{*}$ such that $\phi\left(x_{i}\right)=c$ and $\phi\left(x_{j}\right)=b$. It follows that $t_{(i, j)} \in \operatorname{sol}\left(I^{*}\right)$ if and only if $\left(R,\left(x_{i}, x_{j}\right)\right) \notin C^{*}$, and by definition of $I^{*}$ we have $\left(R,\left(x_{i}, x_{j}\right)\right) \notin C^{*}$ if and only if $t_{(i, j)} \in R_{n}^{*}$. Putting everything together we get that $\operatorname{sol}\left(I^{*}\right) \cap R_{n}=R_{n}^{*}$. Such an instance $I^{*}$ exists for all choices of subrelation $R_{n}^{*}$, so $\Gamma$ shatters $R_{n}$ for every even positive integer $n$. The relation $R_{n}$ contains $(n / 2)^{2}=\Omega\left(n^{2}\right)$ tuples, so the claim follows.

Proof ( of Theorem 13). If $\Gamma$ is rectangular, then by Proposition 16 it has a Mal'tsev derivation. By Proposition 10 we have $\mathrm{CL}_{\Gamma}=O(n)$ and $\Gamma$ is learnable in $O(n \log (n))$ partial membership queries by Theorem 7. Otherwise, by Proposition $17 \mathrm{VC}_{\Gamma}=\Omega\left(n^{2}\right)$ and by Observation 3 any acquisition algorithm for $\Gamma$ must make $\Omega\left(n^{2}\right)$ Boolean queries.

## 6 Beyond binary languages

Theorem 13 shows an interesting dichotomy for binary languages: either $\Gamma$ has a Mal'tsev derivation, or its VCdimension is quadratic. In this section we show that the situation is more complex for ternary languages already as this dichotomy no longer exists (Theorem 18 and Example 3).

Given a domain $D$, we define $p^{*}$ as the ternary partial operation on $D$ with domain $\left(p^{*}\right)=\left\{\left(d_{1}, d_{2}, d_{3}\right) \in\right.$ $\left.D^{3}:\left|\left\{d_{1}, d_{2}, d_{3}\right\}\right| \leq 2\right\}$ and satisfying $p^{*}(x, x, y)=$ $p^{*}(y, x, x)=y$ and $p^{*}(x, y, x)=x$ for all $x, y \in D$. These
three identities are well known and called the Pixley identities; in the terminology of (Lagerkvist and Wahlström 2017) $p^{*}$ would be called the first partial Pixley operation on $D$. The main result of this section is the following.
Theorem 18. Let $\Gamma$ be a ternary language. If $p^{*}$ is a partial polymorphism of $\Gamma$, then $V C_{\Gamma}(n)$ is o $\left(n^{2}\right)$.

We will prove Theorem 18 indirectly, using combinatorial arguments to show that every CSP instance over $\Gamma$ contains at most $o\left(n^{2}\right)$ non-redundant constraints. Observation 4 will then bridge the gap between non-redundancy and VC-dimension. By Observation 5, it suffices to prove Theorem 18 in the case where $\Gamma$ contains a single relation. For the rest of this section we fix a ternary language $\Gamma=\{R\}$ over $D$ which admits $p^{*}$ as a partial polymorphism.

The next lemma establishes a simple but important property that $R$ shares with rectangular relations. As in Section 5, we let $\sim_{1}, \sim_{2}$ and $\sim_{3}$ be three binary relations over $D$ given by:

$$
\begin{aligned}
& d \sim_{1} d^{\prime} \Longleftrightarrow \exists d_{2}, d_{3} \in D \mid\left(d, d_{2}, d_{3}\right),\left(d^{\prime}, d_{2}, d_{3}\right) \in R \\
& d \sim_{2} d^{\prime} \Longleftrightarrow \exists d_{1}, d_{3} \in D \mid\left(d_{1}, d, d_{3}\right),\left(d_{1}, d^{\prime}, d_{3}\right) \in R \\
& d \sim_{3} d^{\prime} \Longleftrightarrow \exists d_{1}, d_{2} \in D \mid\left(d_{1}, d_{2}, d\right),\left(d_{1}, d_{2}, d^{\prime}\right) \in R
\end{aligned}
$$

Lemma 19. $\sim_{1}, \sim_{2}$ and $\sim_{3}$ are equivalence relations over $\left\{d \in D \mid \exists\left(d, d_{2}, d_{3}\right) \in R\right\},\left\{d \in D \mid \exists\left(d_{1}, d, d_{3}\right) \in R\right\}$ and $\left\{d \in D \mid \exists\left(d_{1}, d_{2}, d\right) \in R\right\}$, respectively.

Proof. We prove the lemma for $\sim_{1}$. The cases of $\sim_{2}$ and $\sim_{3}$ are identical. By definition $\sim_{1}$ is reflexive and symmetric, so we need only prove transivity. Let $d, d^{\prime}, d^{\prime \prime}$ be such that $d \sim_{1} d^{\prime}$ and $d^{\prime} \sim_{1} d^{\prime \prime}$. By definition, there exist $d_{1}, d_{2}$ such that $t_{1}=\left(d, d_{1}, d_{2}\right) \in R, t_{2}=\left(d^{\prime}, d_{1}, d_{2}\right) \in R$, and $d_{3}, d_{4}$ such that $t_{3}=\left(d^{\prime}, d_{3}, d_{4}\right) \in R, t_{4}=\left(d^{\prime \prime}, d_{3}, d_{4}\right) \in R$. From the identities of $p^{*}$ we get $p^{*}\left(t_{2}, t_{3}, t_{4}\right)=\left(d^{\prime \prime}, d_{1}, d_{2}\right)$ so $d \sim_{1} d^{\prime \prime}$ and $\sim_{1}$ is transitive.

Let $I=(X, C)$ be a CSP instance over $\Gamma$. For $i \in$ $\{1,2,3\}$ we let $X_{i}=\left\{x_{j}^{i} \mid x_{j} \in X\right\}$. We let $J_{I}$ denote the CSP instance $\left(X_{J}, C_{J}\right)$ over $\Gamma \cup\{=\}$, where $X_{J}=X_{1} \cup$ $X_{2} \cup X_{3}$ and $C_{J}$ contains a constraint $c_{J}=\left(R,\left(x_{i}^{1}, x_{j}^{2}, x_{k}^{3}\right)\right)$ for each $c=\left(R,\left(x_{i}, x_{j}, x_{k}\right)\right) \in C$ as well as two constraints $\left(=,\left(x_{j}^{1}, x_{j}^{2}\right)\right),\left(=,\left(x_{j}^{2}, x_{j}^{3}\right)\right)$ for every $x_{j} \in X$.
Lemma 20. Let c be a constraint in $C$. If $c_{J}$ is redundant in $J_{I}$, then $c$ is redundant in $I$.

Proof. We prove the contrapositive. Suppose that $c$ is not redundant in $I$. Then, there exists an assignment $\phi: X \rightarrow D$ such that $\phi$ violates $c$ but satisfies every $c^{\prime} \in C, c^{\prime} \neq c$. If we define $\phi_{J}: X_{J} \rightarrow D$ such that $\phi_{J}\left(x_{j}^{k}\right)=\phi\left(x_{j}\right)$ for all $x_{j}^{k} \in X_{J}$, then for any $c^{\prime} \in C$ we have $\phi_{J}\left(x_{i}^{1}, x_{j}^{2}, x_{k}^{3}\right)=$ $\phi\left(x_{i}, x_{j}, x_{k}\right)$, so $\phi_{J}$ satisfies $c_{J}^{\prime}$ if and only if $c^{\prime} \neq c$. Since $\phi_{J}$ also satisfies the equalities, $c_{J}$ is not redundant.

By Lemma 20 we can reduce the proof of Theorem 18 to finding a bound on the number of non-redundant constraints in $J_{I}$. To this purpose, we will turn $J_{I}$ into a reduced instance $J_{I}^{*}$, which has the same solution set but fewer constraints; any constraint that does not appear in the reduced instance will be redundant.

We propose the following transformation $T$. If there exists three constraints $c_{1}=\left(R,\left(x_{1}, y_{1}, z_{1}\right)\right), c_{2}=$ $\left(R,\left(x_{1}, y_{2}, z_{2}\right)\right)$ and $c_{3}=\left(R,\left(x_{2}, y_{2}, z_{1}\right)\right)$ (of which at least 2 are distinct) then we remove $c_{2}, c_{3}$ from the instance and instead we add three constraints $\left(\sim_{1},\left(x_{1}, x_{2}\right)\right)$, $\left(\sim_{2},\left(y_{1}, y_{2}\right)\right)$ and $\left(\sim_{3},\left(z_{1}, z_{2}\right)\right)$.
Lemma 21. Applying the transformation $T$ leaves the solution set unchanged.
Proof. Let $c_{1}=\left(R,\left(x_{1}, y_{1}, z_{1}\right)\right), c_{2}=\left(R,\left(x_{1}, y_{2}, z_{2}\right)\right)$ and $c_{3}=\left(R,\left(x_{2}, y_{2}, z_{1}\right)\right)$ be the three constraints on which $T$ will be applied. We first prove that the constraints $\left(\sim_{1},\left(x_{1}, x_{2}\right)\right),\left(\sim_{2},\left(y_{1}, y_{2}\right)\right)$ and $\left(\sim_{3},\left(z_{1}, z_{2}\right)\right)$ are implied by $c_{1}, c_{2}$ and $c_{3}$. Suppose that $\phi$ is a solution to the instance before $T$ is applied. Then, from the constraints $c_{1}, c_{2}, c_{3}$ we deduce that $R$ contains the tuples $t_{1}=\left(\phi\left(x_{1}\right), \phi\left(y_{1}\right), \phi\left(z_{1}\right)\right), t_{2}=\left(\phi\left(x_{1}\right), \phi\left(y_{2}\right), \phi\left(z_{2}\right)\right)$ and $t_{3}=\left(\phi\left(x_{2}\right), \phi\left(y_{2}\right), \phi\left(z_{1}\right)\right)$. From the identities of $p^{*}$ we get $p^{*}\left(t_{1}, t_{2}, t_{3}\right)=\left(\phi\left(x_{2}\right), \phi\left(y_{1}\right), \phi\left(z_{1}\right)\right) \in R$, which imply $x_{1} \sim_{1} x_{2}$ from $t_{1}$ and $y_{1} \sim_{2} y_{2}$ from $t_{3}$. Since $y_{1} \sim_{2} y_{2}$ we deduce from $t_{1}$ that $\left(\phi\left(x_{1}\right), \phi\left(y_{2}\right), \phi\left(z_{1}\right)\right) \in R$, and together with $t_{2}$ we obtain $z_{1} \sim_{3} z_{2}$. Hence, adding the three equivalence constraints does not remove any solution.

Next, we prove that the four constraints $c_{1},\left(\sim_{1},\left(x_{1}, x_{2}\right)\right)$, $\left(\sim_{2},\left(y_{1}, y_{2}\right)\right)$ and $\left(\sim_{3},\left(z_{1}, z_{2}\right)\right)$ together imply $c_{2}$ and $c_{3}$. If $\phi$ is a solution to the instance after $T$ has been applied, then we have $\left(\phi\left(x_{1}\right), \phi\left(y_{1}\right), \phi\left(z_{1}\right)\right) \in R$. Since $\phi\left(y_{1}\right) \sim_{2} \phi\left(y_{2}\right)$ and $\phi\left(z_{1}\right) \sim_{3} \phi\left(z_{2}\right)$, we have that $\left(\phi\left(x_{1}\right), \phi\left(y_{2}\right), \phi\left(z_{2}\right)\right) \in$ $R$ and $c_{2}$ is satisfied. The case of $c_{3}$ is symmetric. Putting everything together we get that replacing the constraints $c_{1}, c_{2}$, $c_{3}$ with $c_{1},\left(\sim_{1},\left(x_{1}, x_{2}\right)\right),\left(\sim_{2},\left(y_{1}, y_{2}\right)\right)$ and $\left(\sim_{3},\left(z_{1}, z_{2}\right)\right)$ leaves the solution set unchanged.

The transformation $T$ may seem odd, because it removes two constraints but adds three new constraints each time. The reason why we are making progress is that the introduced constraints are equivalence relations, among which at most a linear number are non-redundant.

Proof of Theorem 18. We establish the theorem for the case where $\Gamma$ contains a single relation $R$. The result will then follow from Observation 5. Let $I$ be a CSP instance over $\Gamma$ and $J_{I}=\left(X_{J}, C_{J}\right)$ be the associated instance over $\Gamma \cup\{=\}$. We apply the transformation $T$ to $J_{I}$ until fixed point, and then remove all redundant equalities and equivalence constraints. (Note that $T$ removes at least one ternary constraint at each step, so a fixed point will eventually be reached.) The resulting instance $J_{I}^{*}$ contains $O(n)$ equalities and equivalence constraints, and a yet unknown number $k$ of ternary constraints with relation $R$. If we define a graph $G=(V, E)$ such that $V=X_{J}$ and $\{x, y\} \in E$ if and only if there exists a constraint $c=\left(R, S_{c}\right)$ in $J_{I}^{*}$ such that both $x$ and $y$ appear in $S_{c}$, then we have $|E|=3 k$, since no two ternary constraints can share two variables unless $T$ is applicable. Furthermore, every adjacent pair of vertices are contained in a unique triangle (given by the unique constraint $c$ that contains the two variables) unless $T$ is applicable. (If another triangle exists, then its two extra edges belong to distinct constraints $c^{\prime}, c^{\prime \prime}$. Then, $T$ is applicable on $\left\{c, c^{\prime}, c^{\prime \prime}\right\}$ because each variable involved in two constraints must be at
the same position in their respective scopes, by definition of $J_{I}$.) Therefore $G$ is locally linear, and by (Ruzsa and Szemerédi 1978) $G$ has $o\left(n^{2}\right)$ edges. It follows that $k=o\left(n^{2}\right)$ and $J_{I}^{*}$ has $o\left(n^{2}\right)$ constraints.

In order to carry this bound to the number of nonredundant constraints in $J_{I}$, observe that every equivalence relation in $J_{I}^{*}$ is implied by at most three ternary constraints of $J_{I}$. Therefore, if we replace each equivalence constraint in $J_{I}$ with these ternary constraints, we obtain an instance with $o\left(n^{2}\right)$ constraints, all of which belong to $C_{J}$, and whose solution set is by Lemma 21 identical to that of $J_{I}$. It follows that $J_{I}$ has $o\left(n^{2}\right)$ non-redundant constraints, and then by Lemma $20 I$ has $o\left(n^{2}\right)$ non-redundant constraints.

Example 3. Consider the ternary relation

$$
R=\left[\begin{array}{lll}
A & A & B \\
A & C & A \\
B & C & C \\
C & A & C \\
C & B & A
\end{array}\right]
$$

over the domain $D=\{A, B, C\}$. For any three distinct tuples $t_{1}, t_{2}, t_{3} \in R$ there exists $i \in\{1,2,3\}$ such that $t_{1}[i], t_{2}[i]$ and $t_{3}[i]$ are three different values, and hence $\left(t_{1}[i], t_{2}[i], t_{3}[i]\right) \notin$ domain $\left(p^{*}\right)$. Furthermore, for any two (not necessarily distinct) tuples $t_{1}, t_{2} \in R$ we have $p^{*}\left(t_{1}, t_{1}, t_{2}\right)=t_{2}, p^{*}\left(t_{2}, t_{1}, t_{1}\right)=t_{2}$ and $p^{*}\left(t_{1}, t_{2}, t_{1}\right)=$ $t_{1}$. Therefore, $p^{*}$ is a partial polymorphism of $R$ and $\mathrm{VC}_{\{R\}}(n)=o\left(n^{2}\right)$ by Theorem 18.

However, $\{R\}$ does not have a Mal'tsev embedding, even on an infinite domain; it follows from Observation 12 that it does not have a Mal'tsev derivation either. To see this, suppose that $R$ embeds in a relation $\hat{R}$ over a (possibly infinite) domain $\hat{D}$ with a Mal'tsev polymorphism $f$. Then, the operation $g$ given by $g\left(x_{1}, \ldots, x_{5}\right)=$ $f\left(x_{1}, f\left(x_{2}, x_{3}, f\left(x_{1}, x_{2}, x_{3}\right)\right), f\left(x_{5}, x_{4}, f\left(x_{1}, x_{2}, x_{3}\right)\right)\right)$ over $\hat{D}$ is also a polymorphism of $\hat{R}$. Applying $g$ to the five tuples of $R$ (from top to bottom) we obtain ( $B, B, B$ ), which belongs to $\hat{R} \cap D^{3}$ but not to $R$, a contradiction.

As a final remark, we note (without proof, due to space constraints) that the relation $R$ of Example 3 is NP-hard. The proof of Theorem 18 gives a polynomial-time algorithm that transforms any CSP instance over $\{R\}$ into an equivalent one with $o\left(n^{2}\right)$ constraints, or, in other word, a kernel with $o\left(n^{2}\right)$ constraints.

## 7 Conclusion

We have shown that the existence of a Mal'tsev derivation implies linear chain length, and that linear chain length implies efficient learnability with partial membership queries. For binary languages, by Theorem 13 these properties are equivalent. Could this also be true in the general case? Does there exist languages with linear chain length but no Mal'tsev derivation, or languages that are efficiently learnable despite having superlinear chain length?

## References

Angluin, D. 1988. Queries and concept learning. Machine Learning 2(4):319-342.
Barto, L.; Krokhin, A. A.; and Willard, R. 2017. Polymorphisms, and how to use them. In The Constraint Satisfaction Problem: Complexity and Approximability. 1-44.
Beldiceanu, N., and Simonis, H. 2012. A model seeker: Extracting global constraint models from positive examples. In Proceedings of the Eighteenth International Conference on Principles and Practice of Constraint Programming (CP'12), 141-157.
Bessiere, C.; Coletta, R.; Koriche, F.; and O'Sullivan, B. 2005. A SAT-based version space algorithm for acquiring constraint satisfaction problems. In Proceedings of the Sixteenth European Conference on Machine Learning (ECML'05), 23-34.
Bessiere, C.; Coletta, R.; Hebrard, E.; Katsirelos, G.; Lazaar, N.; Narodytska, N.; Quimper, C.; and Walsh, T. 2013. Constraint acquisition via partial queries. In Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI'13), 475-481.
Bessiere, C.; Koriche, F.; Lazaar, N.; and O’Sullivan, B. 2017. Constraint acquisition. Artif. Intell. 244:315-342.

Bulatov, A. A., and Dalmau, V. 2006. A simple algorithm for mal'tsev constraints. SIAM Journal on Computing 36(1):1627.

Couceiro, M.; Haddad, L.; and Lagerkvist, V. 2019. Fine-grained complexity of constraint satisfaction problems through partial polymorphisms: A survey. In Proceedings of the Forty-Ninth International Symposium on MultipleValued Logic (ISMVL'19), 170-175.
Hooker, J. N., and van Hoeve, W. J. 2018. Constraint programming and operations research. Constraints 23(2):172195.

Idziak, P. M.; Markovic, P.; McKenzie, R.; Valeriote, M.; and Willard, R. 2010. Tractability and learnability arising from algebras with few subpowers. SIAM J. Comput. 39(7):3023-3037.
Jonsson, P.; Lagerkvist, V.; and Roy, B. 2017. Time complexity of constraint satisfaction via universal algebra. In Proceedings of the Forty-Second International Symposium on Mathematical Foundations of Computer Science (MFCS'17), 17:1-17:15.
Kazda, A. 2011. Maltsev digraphs have a majority polymorphism. European Journal of Combinatorics 32(3):390-397. Lagerkvist, V., and Wahlström, M. 2017. Kernelization of constraint satisfaction problems: A study through universal algebra. In Proceedings of the Twenty-Third International Conference on Principles and Practice of Constraint Programming (CP'17), 157-171.
Lagerkvist, V., and Wahlström, M. 2017. Kernelization of constraint satisfaction problems: A study through universal algebra. arXiv:1706.05941, rev. 1.
Négrevergne, B., and Guns, T. 2015. Constraint-based sequence mining using constraint programming. In Pro-
ceedings of the Twelfth International Conference on Integration of AI and OR Techniques in Constraint Programming (CPAIOR'15), 288-305.
Ruzsa, I. Z., and Szemerédi, E. 1978. Triple systems with no six points carrying three triangles. Combinatorics 18:939945.

Shaw, P. 1998. Using constraint programming and local search methods to solve vehicle routing problems. In Proceedings of the Fourth International Conference on Principles and Practice of Constraint Programming (CP'98), 417431.
van Beek, P., and Chen, X. 1999. Cplan: A constraint programming approach to planning. In Proceedings of the Sixteenth National Conference on Artificial Intelligence and Eleventh Conference on Innovative Applications of Artificial Intelligence (AAAI-IAAI'99), 585-590.


[^0]:    Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

