Formal languages and the model theory of finitely generated free monoids

Christa Hawthorne*

August 2016

Abstract

The failure of model-theoretic tameness properties such as NIP and categoricity is established for finitely generated free monoids. A characterization is given of the quantifier-free definable sets in one variable. These are shown to be star-free. Working in the expansion (M, A) of a finitely generated free monoid by a predicate for a distinguished set, it is shown that A is a regular language precisely when (M, A) is of φ-rank zero for a particular choice of φ. The latter observation is generalized to arbitrary monoids with recognizable subsets replacing regular languages.

Contents

1 Introduction

The fields of formal languages and automata theory in computer science are known to have a rich interplay with logic; see for example [1] and [4, Chapter XVIII]. One of the more fruitful connections involves characterizing classes of formal languages by treating strings as structures in a certain signature, and asking how much logical power is needed to axiomatize a given language. In this paper, we take a different approach: we consider languages as subsets of finitely generated free monoids, and hope to derive results about the former by studying model-theoretic properties of the latter. We assume no knowledge of formal languages or automata theory; all of the requisite material will be covered as needed.

We now describe our results. We first establish the following basic properties:

*This work was partially supported by an NSERC CGS-M.

- Any non-empty class of (non-trivial) free monoids is not elementary.
- The complete theory of a finitely generated free monoid is not model-complete, not $\aleph_0$-categorical (or even small), and not stable; if the number of generators is at least two, it has the independence property.
- A finitely generated free monoid is a prime model of its theory and is strongly ω-homogeneous, though it is not ω-saturated.
- In contrast to the model theory of free groups, non-isomorphic finitely generated free monoids are not elementarily equivalent: the number of generators is captured in the first-order theory.

Next we study the quantifier-free definable sets in one variable. We give an algebraic characterization thereof, showing in particular that in the formal languages hierarchy they land strictly between the finite languages and the star-free languages. On the other hand, we do not believe that the definable sets fit well into this hierarchy: we give an example of a definable set in one variable that is not regular, and we give an example of a regular language that we conjecture is not definable.

In the wake of the above conjecture that not all regular languages are definable, we consider the expansion of a finitely generated free monoid by a unary predicate symbol for an identified language; we consider the implications that regularity of the identified language has on the expanded structure. More generally, we consider the expansion of an arbitrary monoid by a unary predicate for an identified subset, and consider the implications of the identified set being *recognizable*—this is a generalization of regularity to the setting of an arbitrary monoid. We show that A is a recognizable subset of a monoid M if and only if the theory of (M, A) has φ-rank zero, where $\varphi(x; u, v)$ is $uxv \in A$. An exposition of local types and φ-rank is included.

I warmly thank my advisor, Rahim Moosa, for excellent guidance, thorough editing, and many helpful discussions.

2 Preliminaries from the study of formal languages

We begin by introducing the basic concepts we will need from the study of formal languages.

2.1 Basic notions

Throughout this paper, we will assume Σ is a finite non-empty set, called an *alphabet.* Common choices will be $n < \omega$ or $\{\mathrm{a}, \mathrm{b}, \ldots, \mathrm{z}\}$.

Definition 2.1. A *string* (or *word*) over Σ is a finite (possibly empty) tuple of elements of Σ. We use Σ^* to denote the set of strings over Σ; formally, we have

$$\Sigma^* = \bigcup_{n<\omega} \Sigma^n$$

where Σ^n denotes the n^{th} Cartesian power of Σ. We use ε to denote the empty word; we let $\Sigma^+ = \Sigma^* \setminus \{\varepsilon\}$ be the set of non-empty words over Σ. Given a string $a \in \Sigma^*$ we denote its length by $|a|$; given strings $a, b \in \Sigma^*$ we denote their concatenation by $a \cdot b$ or simply ab.

Remark 2.2. Concatenation is associative and has ε as an identity; hence $(\Sigma^*, \varepsilon, \cdot)$ forms a monoid. In fact, it is the free monoid on Σ; one can see this by verifying the universal property directly.

Definition 2.3. Suppose $a, b \in \Sigma^*$.

- We say b is a *substring* of a if there are $c_1, c_2 \in \Sigma^*$ such that $a = c_1 b c_2$.
- We say b is a *prefix* of a if there is $c \in \Sigma^*$ such that $a = bc$.
- We say b is a *suffix* of a if there is $c \in \Sigma^*$ such that $a = cb$.
- We say a and b are *conjugate* if there are $c_1, c_2 \in \Sigma^*$ such that $a = c_1 c_2$ and $b = c_2 c_1$. (Roughly speaking, if b is a cyclic shift of a.)

We say b is a *proper substring, prefix, or suffix* of a if b is a substring, prefix, or suffix of a, respectively, and $b \neq a$.

Example 2.4. In $\Sigma = \{ \mathrm{a}, \mathrm{b}, \dots, \mathrm{z} \}$, let a be "bandersnatch". Then:

- "band" is a prefix of a.
- "snatch" is a suffix of a.
- "ersna" is a substring of a.
- "snatchbander" is a conjugate of a. (Via $c_1 =$ "bander" and $c_2 =$ "snatch".)

Remark 2.5. The notion of conjugacy, as it is phrased above, makes sense in any monoid, and in particular in any group. Indeed, one can verify that for groups this notion coincides with the usual notion of conjugacy in group theory.

Remark 2.6. If $a, b \in \Sigma^*$ are conjugate, then $|a| = |b|$.

2.2 Some facts about words

We note two theorems of Lyndon and Schützenberger; proofs can be found in [7, Section 2.3]. (Schützenberger was French; his name is pronounced accordingly.)

The first theorem of Lyndon-Schützenberger concerns when a word can be both a prefix and a suffix of another word:

Theorem 2.7 ([7, Theorem 2.3.2])**.** *Suppose $a, b, c \in \Sigma^+$. Then $ab = bc$ if and only if there are $d_1, d_2 \in \Sigma^*$ and $n < \omega$ such that*

1. $a = d_1 d_2$
2. $c = d_2 d_1$
3. $b = (d_1 d_2)^n d_1$

Note that in this case we have

$$b = a^n d_1 = (d_1 d_2)^n d_1 = d_1 (d_2 d_1)^n = d_1 c^n$$

Example 2.8. Let $\Sigma = \{ 0, 1 \}$; let

$$\begin{aligned} a &= 011 \\ b &= 0110 \\ c &= 110 \end{aligned}$$

Then $ab = 0110110 = bc$, and if we take $d_1 = 0$, $d_2 = 11$, and $n = 1$, we see that Theorem 2.7 indeed holds in this case.

The second theorem of Lyndon-Schützenberger concerns when two words commute:

Theorem 2.9 ([7, Theorem 2.3.3])**.** *Suppose $a, b \in \Sigma^+$. Then $ab = ba$ if and only if there is $c \in \Sigma^*$ and $n, m < \omega$ such that $a = c^n$ and $b = c^m$.*

We now turn to a unique factorization result on words; proofs and more information can again be found in [7, Section 2.3].

Definition 2.10. We say $a \in \Sigma^+$ is *primitive* if there does not exist $b \in \Sigma^*$ and $1 < n < \omega$ such that $a = b^n$.

Example 2.11. In $\Sigma = \{ \mathrm{a}, \dots, \mathrm{z} \}$ we have that "jubjub" = ("jub")^2 is not primitive, whereas "jubjubbird" is primitive.

Remark 2.12. If $|a|$ is prime and a does not take that form ℓ^n for some $\ell \in \Sigma$ and $n < \omega$, then a is primitive.

Remark 2.13. If $a, b \in \Sigma^+$ and a is primitive, we can strengthen the second theorem of Lyndon-Schützenberger (Theorem 2.9) to the following: $ab = ba$ if and only if b is a power of a.

Theorem 2.14 ([7, Theorem 2.3.4])**.** *Suppose $a \in \Sigma^+$. Then there is a unique primitive $b \in \Sigma^+$ and $0 < n < \omega$ such that $a = b^n$.*

Definition 2.15. Suppose $a \in \Sigma^+$; write $a = b^n$ where $b \in \Sigma^+$ is primitive and $0 < n < \omega$. We define $\sqrt{a} = b$ and $\deg(a) = n$; thus $a = (\sqrt{a})^{\deg(a)}$.

Example 2.16. Let $\Sigma = \{\, \mathrm{a}, \dots, \mathrm{z} \,\}$; let $a =$ "tumtum". Then $\sqrt{a} =$ "tum", and $\deg(a) = 2$.

Remark 2.17. Comparing lengths in the equation $a = (\sqrt{a})^{\deg(a)}$, we find that $|a| = |\sqrt{a}| \deg(a)$. One also notices that $\sqrt{a}$ is both a prefix and a suffix of a.

Remark 2.18. a is primitive if and only if $\deg(a) = 1$.

A useful characterization of primitivity:

Proposition 2.19. *Suppose $a \in \Sigma^+$. Then a is primitive if and only if there do not exist $b_1, b_2 \in \Sigma^+$ such that $a = b_1 b_2 = b_2 b_1$.*

Proof.

($\implies$) Suppose we have $b_1, b_2 \in \Sigma^+$ such that $a = b_1 b_2 = b_2 b_1$; we will show that a is not primitive. Now, since $b_1 b_2 = b_2 b_1$, the second theorem of Lyndon-Schützenberger yields that there is $c \in \Sigma^*$ and $n, m < \omega$ such that $b_1 = c^n$ and $b_2 = c^m$. In particular, since b_1 and b_2 are non-empty, we find that n and m are non-zero; hence $n + m > 1$. But now $a = b_1 b_2 = c^n c^m = c^{n+m}$; so a is not primitive.

($\impliedby$) Suppose a is not primitive; say $a = c^n$ where $c \in \Sigma^*$ and $1 < n < \omega$. Since a is non-empty, so too is c; since $n - 1 > 0$, we thus get that c^{n-1} is non-empty. Now, letting $b_1 = c$ and $b_2 = c^{n-1}$, we find that b_1 and b_2 are both non-empty, and $a = c^n = b_1 b_2 = b_2 b_1$. ☐ Proposition 2.19

Corollary 2.20. *Suppose $a \in \Sigma^+$ is primitive and $b \in \Sigma^*$ is a conjugate of a. Then there are unique $c_1, c_2 \in \Sigma^+$ such that $a = c_1 c_2$ and $b = c_2 c_1$.*

Proof. Suppose we have $c_1, c_2, c_1', c_2' \in \Sigma^+$ such that $a = c_1 c_2 = c_1' c_2'$ and $b = c_2 c_1 = c_2' c_1'$; without loss of generality we assume that $|c_2'| \geq |c_2|$. Hence c_2 is a prefix of c_2', and there is $d \in \Sigma^*$ such that $c_2' = c_2 d$. Then since $c_2 c_1 = b = c_2' c_1' = c_2 d c_1'$, we may cancel c_2 to get that $c_1 = d c_1'$. Substituting into $a = c_1 c_2 = c_1' c_2'$, we find that $a = d c_1' c_2 = c_1' c_2 d$. But $c_1' c_2 \neq \varepsilon$, and a is primitive; so, by Proposition 2.19, we get that $d = \varepsilon$. So $c_2' = c_2$, and thus $c_1' = c_1$; so there are unique such $c_1, c_2 \in \Sigma^+$, as desired. ☐ Corollary 2.20

2.3 Languages, operations on languages, and classes of languages

We now discuss languages; a good reference for this subsection is [7, Chapter 1].

Definition 2.21. A *language* is any subset of Σ^*.

Example 2.22. If $\Sigma = \{\, \mathrm{a}, \mathrm{b}, \dots, \mathrm{z} \,\}$, one might consider the language of words that appear in the Oxford English dictionary.

We now define some operations on languages:

Definition 2.23. Suppose $A, B \subseteq \Sigma^*$ are languages. We define:

- the *concatenation* of A and B to be $AB = A \cdot B = \{\, ab : a \in A, b \in B \,\}$.
- the *Kleene star* of A to be

$$A^* = \bigcup_{n<\omega} A^n = \{\, a_1 \cdots a_n : a_1, \dots, a_n \in A, n < \omega \,\}$$

("Kleene" is pronounced to rhyme with "zany".) As a notational convenience, we let

$$A^+ = A^* \setminus \{\, \varepsilon \,\} = \bigcup_{0<n<\omega} A^n$$

Example 2.24. Let $\Sigma = \{ \text{a}, \ldots, \text{z} \}$; let $A = \{ \text{“uff”}, \text{“beam”} \}$ and $B = \{ \text{“ish”} \}$. Then $AB = \{ \text{“uffish”}, \text{“beamish”} \}$, and

$$A^* = \{ \varepsilon, \text{“uff”}, \text{“beam”}, \text{“uffuff”}, \text{“uffbeam”}, \text{“beamuff”}, \text{“beambeam”}, \text{“uffuffuff”}, \ldots \}$$

Remark 2.25. Note that an element of AB may not factor uniquely as ab for $a \in A$ and $b \in B$. Consider for example $\Sigma = \{ \text{a}, \ldots, \text{z} \}$, $A = \{ \text{“snicker”}, \text{“snickers”} \}$, and $B = \{ \text{“nack”}, \text{“snack”} \}$. Then

$$\text{“snickersnack”} = (\text{“snicker”})(\text{“snack”}) = (\text{“snickers”})(\text{“nack”}) \in AB$$

When combining operations, it is understood that Kleene star and complementation bind tighter than concatenation, which binds tighter than union and intersection. For example, the correct parenthesization of $AB^cC^* \cup D$ would be $(A(B^c)(C^*)) \cup D$. (Note that concatenation of languages is associative, so we are quite happy to leave $A(B^c)(C^*)$ without further parentheses.)

When convenient, we will allow a word a to stand in for $\{ a \}$ when expressing a language using the above operations; for example, when we write $a^*b \cup c^*d$ it will be understood to mean $\{ a \}^*\{ b \} \cup \{ c \}^*\{ d \}$.

We now come to an important class of languages, both in the theory of formal languages and in applications:

Definition 2.26. The class of *regular languages* is the smallest class of languages containing the finite languages and closed under finite union, concatenation, and Kleene star; i.e.

- Every finite language is regular.
- If A and B are regular, then so are $A \cup B$, AB, and A^*.
- Nothing else is.

Example 2.27. Let $\Sigma = \{ \text{a}, \ldots, \text{z} \}$. Let $A \subseteq \Sigma^*$ be the set of words that consist of a “br” followed by a vowel followed by an even number of “l” followed by another vowel followed by a “g”. Then A is regular; indeed,

$$A = \text{“br”}(\{ \text{“a”}, \text{“e”}, \text{“i”}, \text{“o”}, \text{“u”} \})(\text{“ll”})^*(\{ \text{“a”}, \text{“e”}, \text{“i”}, \text{“o”}, \text{“u”} \})\text{“g”}$$

The following result follows from an automata-theoretic characterization of regular languages that we won’t touch on here.

Proposition 2.28 ([7, Corollary 1.4.3])**.** *If $A \subseteq \Sigma^*$ is regular, then so is $A^c = \Sigma^* \setminus A$.*

Since the union of regular languages is regular, it follows that the class of regular languages is closed under Boolean combinations.

Regular languages and finite languages are two entries in a hierarchy of classes of languages, ordered by complexity. We mention in passing one other class in this hierarchy: the class of *context-free languages*. While a proper definition would take us too far afield, we will mention that every regular language is context-free, and that the class of context-free languages contains significantly more than the class of regular languages.

3 First model-theoretic properties of finitely generated free monoids

We now begin to study $(\Sigma^*, \varepsilon, \cdot)$ from a model-theoretic perspective; that is, we work in the language $L_{\text{Mon}} = \{ \varepsilon, \cdot \}$, where

- ε is a constant symbol
- $\cdot$ is a binary function symbol

and we view Σ^* as an L_{Mon}-structure in the natural way.

We start with some elementary properties that can be seen without delving too deeply.

Remark 3.1. Consider the L_{Mon}-formula $\lambda_1(x)$ given by

$$(x \neq \varepsilon) \wedge \neg(\exists u \exists v (u \neq \varepsilon \wedge v \neq \varepsilon \wedge x = uv))$$

Then $\lambda_1(\Sigma^*)$ consists of the non-empty words that cannot be written as the concatenation of two non-empty words; it is easily seen that these are exactly the elements of Σ. So $\lambda_1(\Sigma^*) = \Sigma$, and Σ is $\emptyset$-definable in $(\Sigma^*, \varepsilon, \cdot)$. Consider now the L_{Mon}-formulae

$$\begin{aligned}
\lambda_0(x) &= (x = \varepsilon) \\
\lambda_n(x) &= \exists u_1 \cdots \exists u_n \left(x = u_1 \cdots u_n \wedge \bigwedge_{i=1}^{n} \lambda_1(u_i) \right) \text{ (for } n \geq 2)
\end{aligned}$$

Then $\lambda_n(\Sigma^*)$ is the set of words of length n, which is thus $\emptyset$-definable in $(\Sigma^*, \varepsilon, \cdot)$.

This quickly distinguishes the model theory of free monoids from the model theory of free groups. Whereas [6] shows that any two free groups on at least two generators are elementarily equivalent, here we have the following:

Corollary 3.2. *Suppose $n < m < \omega$ and κ is an infinite cardinal. Then the free monoids on n, m, and κ generators are pairwise elementarily inequivalent.*

Proof. Let $\mathcal{M}_n$, $\mathcal{M}_m$, and $\mathcal{M}_\kappa$ be the free monoids on n, m, and κ generators, respectively. Note that

$$\begin{aligned}
\mathcal{M}_n &\models \exists^{=n} x(\lambda_1(x)) \wedge \neg\exists^{=m} x(\lambda_1(x)) \\
\mathcal{M}_m &\models \neg\exists^{=n} x(\lambda_1(x)) \wedge \exists^{=m} x(\lambda_1(x)) \\
\mathcal{M}_\kappa &\models \neg\exists^{=n} x(\lambda_1(x)) \wedge \neg\exists^{=m} x(\lambda_1(x))
\end{aligned}$$

Hence $\mathcal{M}_n$, $\mathcal{M}_m$, and $\mathcal{M}_\kappa$ are pairwise elementarily inequivalent, as desired. ☐ Corollary 3.2

Another consequence of Remark 3.1 is that first-order logic fails badly to capture when a monoid is free.

Corollary 3.3. *Any non-trivial free monoid has an elementary extension that is not free.*

In particular, no non-empty class of non-trivial free monoids is elementary.

Proof. Suppose $\mathcal{M}$ is the free monoid on $X \neq \emptyset$. (Note that we do not require X to be finite.) Consider

$$\Xi(x) = \{ \neg\lambda_n(x) : n < \omega \}$$

Observe that none of the analysis in Remark 3.1 relied on the alphabet being finite; we may thus freely apply the conclusions of Remark 3.1 to $\mathcal{M}$. Hence if $a \in X$ and $m < \omega$, then $\mathcal{M} \models \neg\lambda_n(a^m)$ for all $n < m$; hence Ξ is finitely satisfiable in $\mathcal{M}$. By compactness, there is $\mathcal{N} \succeq \mathcal{M}$ with a realization $a \in N$ of Ξ. But Remark 3.1 yields that Ξ cannot be realized in any free monoid; so $\mathcal{N}$ is not free. ☐ Corollary 3.3

Remark 3.1 also yields the following:

Proposition 3.4. $(\Sigma^*, \varepsilon, \cdot)$ *is a prime model of its theory.*

Proof. Since we are in a countable language, we have that a model is prime if and only if it is countable and atomic. Since Σ is finite, we get that $(\Sigma^*, \varepsilon, \cdot)$ is countable; it remains to check that it is atomic. Suppose then that $a_1, \ldots, a_n \in \Sigma^*$; we wish to check that $p(x_1, \ldots, x_n) = \mathrm{tp}(a_1 \cdots a_n/\emptyset)$ is isolated. Consider

$$\varphi(x_1, \ldots, x_n) = \bigwedge_{i=1}^{n} \lambda_{|a_i|}(x_i)$$

Claim 3.5. *φ is algebraic.*

Proof. Note that for any $n < \omega$ there are exactly $|\Sigma|^n$ words of length n; hence $|\lambda_n(\Sigma^*)| = |\Sigma|^n$. But

$$\varphi(\Sigma^*) = \{ (b_1, \ldots, b_n) \in (\Sigma^*)^n : b_1 \in \lambda_{|a_1|}, \ldots, b_n \in \lambda_{|a_N|} \} = \prod_{i=1}^{n} \lambda_{|a_i|}(\Sigma^*)$$

Hence

$$|\varphi(\Sigma^*)| = \left| \prod_{i=1}^{n} \lambda_{|a_i|}(\Sigma^*) \right| = \prod_{i=1}^{n} \left| \lambda_{|a_i|}(\Sigma^*) \right| = \prod_{i=1}^{n} |\Sigma|^{|a_i|} < \aleph_0$$

So φ is algebraic, as desired. ☐ Claim 3.5

But $\varphi \in p$. So p is algebraic, and is thus isolated.

So every realized type is isolated. So $(\Sigma^*, \varepsilon, \cdot)$ is atomic. So $(\Sigma^*, \varepsilon, \cdot)$ is a prime model of its theory. $\square$ Proposition 3.4

Corollary 3.6. *$(\Sigma^*, \varepsilon, \cdot)$ is ω-homogeneous; in particular, since it is countable, it is strongly ω-homogeneous.*

Unfortunately, these seem to be the only nice model-theoretic properties satisfied by $(\Sigma^*, \varepsilon, \cdot)$. The rest of this section will be devoted to cataloguing the failures of other properties one might have hoped for.

Remark 3.7. $\{\neg\lambda_n(x) : n < \omega\}$ is a partial type over $\emptyset$ that is not realized in Σ^*; hence $(\Sigma^*, \varepsilon, \cdot)$ is not ω-saturated.

Proposition 3.8. *$\mathrm{Th}(\Sigma^*, \varepsilon, \cdot)$ is not small.*

Proof. We check that $|S_1(\mathrm{Th}(\Sigma^*, \varepsilon, \cdot))| > \aleph_0$. For $p \in \mathbb{N}$ prime, let $\varphi_p(x)$ be $\exists y(x = y^p)$. Given a set A of primes, let

$$\Xi_A(x) = \{\varphi_p : p \in A\} \cup \{\neg\varphi_p : p \notin A\}$$

Claim 3.9. *Ξ_A is a partial type for any set of primes A.*

Proof. We show that Ξ_A is finitely satisfiable in $(\Sigma^*, \varepsilon, \cdot)$. Suppose we are given a finite subset of Ξ_A, say

$$\{\varphi_p : p \in B\} \cup \{\neg\varphi_p : p \in C\}$$

where $B \subseteq A$ and $C \subseteq \mathbb{N} \setminus A$ are finite sets of primes. Pick some $\ell \in \Sigma$; let

$$N = \prod_{p \in B} p$$

and let $a = \ell^N$. Then for any prime p we have $(\Sigma^*, \varepsilon, \cdot) \models \varphi_p(a)$ if and only if $p \mid N$; i.e. if and only if $p \in B$. In particular, we get that $(\Sigma^*, \varepsilon, \cdot) \models \varphi_p(a)$ for $p \in B$ and $(\Sigma^*, \varepsilon, \cdot) \models \neg\varphi_p(a)$ for $p \in C$.

So Ξ_A is finitely consistent, and is thus a partial type. $\square$ Claim 3.9

For each set A of primes, let p_A be any extension of Ξ_A to a complete type; then, since the Ξ_A are mutually inconsistent, we get that the p_A are distinct elements of $S_1(T)$. But there are $2^{\aleph_0}$-many sets of primes; so $|S_1(T)| \geq 2^{\aleph_0} > \aleph_0$.

So $\mathrm{Th}(\Sigma^*, \varepsilon, \cdot)$ is not small. $\square$ Proposition 3.8

Corollary 3.10. *$\mathrm{Th}(\Sigma^*, \varepsilon, \cdot)$ is neither $\aleph_0$-categorical nor ω-stable.*

Proof. Failure of $\aleph_0$-categoricity follows from Ryll-Nardzewski, since we are working in a countable language; failure of ω-stability follows from the definitions. $\square$ Corollary 3.10

Proposition 3.11. *$\mathrm{Th}(\Sigma^*, \varepsilon, \cdot)$ is not stable.*

Proof. Let $\varphi(x, y)$ be $(x \neq y) \wedge \exists z(x = yz)$; that is, $\varphi(x, y)$ asserts that y is a proper prefix of x. Fix $\ell \in \Sigma$; for $i < \omega$, let $a_i = b_i = \ell^i$. Then

$$(\Sigma^*, \varepsilon, \cdot) \models \varphi(a_i, b_j) \iff i < j$$

and φ has the order property. So φ is not stable; so $\mathrm{Th}(\Sigma^*, \varepsilon, \cdot)$ is not stable. $\square$ Proposition 3.11

In fact, the situation is worse:

Proposition 3.12. *$\mathrm{Th}(\Sigma^*, \varepsilon, \cdot)$ has the independence property if $|\Sigma| \geq 2$.*

Proof. Let $\varphi(x, y) = \exists u \exists v(x = uyv)$; that is, $\varphi(x, y)$ asserts that y is a substring of x. I claim that φ has the independence property.

Fix distinct $\ell_1, \ell_2 \in \Sigma$; for $i < \omega$, let $a_i = \ell_1 \ell_2^{i+1} \ell_1$. Suppose $A \subseteq \omega$; I claim that

$$\Xi(x) = \{\varphi(x, a_i) : i \in A\} \cup \{\neg\varphi(x, a_i) : i \notin A\}$$

is consistent. We will show that it is finitely consistent; suppose then that we are given a finite subset, say of the form

$$\Phi(x) = \{ \varphi(x, a_i) : i \in B \} \cup \{ \neg\varphi(x, a_i) : i \in C \}$$

where $B \subseteq A$ and $C \subseteq \omega \setminus A$ are finite. Let $b \in \Sigma^*$ be the concatenation of $(\ell_1 \ell_2^{i+1} \ell_1 : i < \omega, i \in B)$. It is then easily seen that given $i < \omega$ we have that $\ell_1 \ell_2^{i+1} \ell_1$ is a substring of b if and only if $i \in B$; hence b is a realization of Φ.

So Ξ is consistent; so φ has the independence property, and $\mathrm{Th}(\Sigma^*, \varepsilon, \cdot)$ has the independence property. $\square$ Proposition 3.12

Note that if $|\Sigma| = 1$ then $(\Sigma^*, \varepsilon, \cdot) \cong (\mathbb{N}, 0, +)$ is just Presburger arithmetic, which is NIP; see [8, Section A.2].

Proposition 3.13. *$\mathrm{Th}(\Sigma^*, \varepsilon, \cdot)$ is not model-complete.*

Proof. Let $f\colon (\Sigma^*, \varepsilon, \cdot) \to (\Sigma^*, \varepsilon, \cdot)$ be given by $\ell \mapsto \ell\ell$ for $\ell \in \Sigma$; since Σ^* is the free monoid on Σ, this specifies a unique homomorphism of monoids. It is clear that f is injective. Since L_{Mon} contains no relation symbols, we then get that f is an L_{Mon}-embedding. But f is not an elementary embedding: for $\ell \in \Sigma$ we have

$$(\Sigma^*, \varepsilon, \cdot) \not\models \exists x(\ell = x^2)$$
$$(\Sigma^*, \varepsilon, \cdot) \models \exists x(f(\ell) = x^2)$$

So we have an embedding of models of $\mathrm{Th}(\Sigma^*, \varepsilon, \cdot)$ that is not elementary; hence $\mathrm{Th}(\Sigma^*, \varepsilon, \cdot)$ is not model-complete. $\square$ Proposition 3.13

We immediately get the following:

Corollary 3.14. *$\mathrm{Th}(\Sigma^*, \varepsilon, \cdot)$ does not admit quantifier elimination.*

In fact, the proof of Proposition 3.13 actually gives us an example of a formula that is not equivalent to a quantifier-free one, namely $\exists x(y = x^2)$.

4 Definable sets

We exhibit some concepts from formal languages that are definable in $(\Sigma^*, \varepsilon, \cdot)$.

Proposition 4.1.

1. *The relation "is a conjugate of" is $\emptyset$-definable.*
2. *Primitivity is $\emptyset$-definable.*
3. *If $a \in \Sigma^*$ is primitive, then a^* is quantifier-free definable in a uniform way.*
4. *If $a \in \Sigma^+$, then a^* is definable; however, the definition is not uniform in a.*
5. *Suppose $A, B \subseteq \Sigma^*$ are definable. Then AB is definable.*

Proof.

1. We can take $\varphi(x, y)$ to be
$$\exists u \exists v((x = uv) \wedge (y = vu))$$
2. Using Proposition 2.19, we can take $\varphi(x)$ to be
$$\neg\exists u \exists v((u \neq \varepsilon) \wedge (v \neq \varepsilon) \wedge (x = uv = vu))$$
3. Using Remark 2.13, we can take $\varphi(x)$ to be $ax = xa$.

4. Take $\varphi(x)$ to be $(ax = xa) \wedge (\exists y(x = y^{\deg(a)}))$. Using Theorem 2.9, it is routine to check that $\varphi(\Sigma^*) \supseteq a^*$; it remains to check that $\varphi(\Sigma^*) \subseteq a^*$.

   Suppose then that $b \in \varphi(\Sigma^*)$; that is, suppose $ab = ba$ and we have $c \in \Sigma^*$ such that $b = c^{\deg(a)}$. By Theorem 2.9, we get that a and b are both powers of some $d \in \Sigma^*$; hence we can write both a and b as powers of $\sqrt{d}$, and we thus see that $\sqrt{b} = \sqrt{d} = \sqrt{a}$. Now, since $b = c^{\deg(a)}$, we can write b as a power of $\sqrt{c}$; hence $\sqrt{c} = \sqrt{b} = \sqrt{a}$, and $c = (\sqrt{a})^n$ for some $n < \omega$. But then $b = c^{\deg(a)} = (\sqrt{a})^{n \deg(a)} = a^n \in a^*$, as desired.

5. Let A and B be defined by $\varphi_A(x)$ and $\varphi_B(x)$, respectively. We can then take $\varphi(x)$ to be

$$\exists y \exists z((x = yz) \wedge \varphi_A(y) \wedge \varphi_B(z))$$

□ Proposition 4.1

Example 4.2. $\{ a^2 : a \in \Sigma^* \}$ is definable: we can take $\varphi(x)$ to be $\exists y(x = y^2)$. In fact, $\{ a^2 : a \in \Sigma^* \}$ is known not to be regular. (It isn't even context-free; see [9, Example 2.38].) Hence we have an example of a language that is definable but not regular.

Ideally, we would have some examples of languages that aren't definable; however, checking that something isn't definable is troublesome. (They must of course exist, since our language and domain set are countable.) The best we have is a conjecture:

Conjecture 4.3. $(\Sigma\Sigma)^* = \{ a \in \Sigma^* : |a| \in 2\mathbb{N} \}$ is not definable.

If this conjecture holds then we have an example of a regular language that isn't definable; this, combined with Example 4.2, would show that the class of definable languages doesn't fit well into the formal languages hierarchy.

A weaker conjecture would be the following:

Conjecture 4.4. $\{ (a, b) \in (\Sigma^*)^2 : |a| = |b| \}$ is not definable.

4.1 A characterization of the quantifier-free definable subsets of Σ^*

Our basic examples of quantifier-free definable subsets of Σ^* will be those defined by formulae of the form $ax = xb$ for $a, b \in \Sigma^+$; in fact, these will turn out to be the building blocks for all quantifier-free subsets of Σ^*. The first theorem of Lyndon-Schützenberger (Theorem 2.7) gives us a characterization of the x such that $ax = xb$; however, for our purposes, a variation on their result will be more helpful. First, however, we present some results on conjugates:

Lemma 4.5. *Suppose $a, b \in \Sigma^+$ are conjugate; say $a = c_1c_2$ and $b = c_2c_1$. Write $c_1 = (\sqrt{a})^n c_1'$ and $c_2 = c_2'(\sqrt{a})^m$ where $|c_1'| < |\sqrt{a}|$, $|c_2'| < |\sqrt{a}|$, and $m + n + 1 = \deg(a)$. Then $\sqrt{a} = c_1'c_2'$ and $\sqrt{b} = c_2'c_1'$.*

Roughly speaking, the lemma asserts that two words are conjugate only if their roots are, and the way in which they are conjugate arises from the way in which their roots are conjugate. Note that since $\sqrt{a}$ and $\sqrt{b}$ are primitive, Corollary 2.20 yields that c_1' and c_2' are unique, if they exist.

Proof. Without loss of generality, we may assume $|\sqrt{a}| \leq |\sqrt{b}|$.

Since $(\sqrt{a})^{\deg(a)} = a = c_1c_2 = (\sqrt{a})^n c_1'c_2'(\sqrt{a})^m$ and $m + n + 1 = \deg(a)$, we immediately get that $\sqrt{a} = c_1'c_2'$. But note also that

$$\begin{aligned} b &= c_2c_1 \\ &= (c_2'(\sqrt{a})^m)((\sqrt{a})^n c_1') \\ &= c_2'(\sqrt{a})^{\deg(a)-1} c_1' \\ &= c_2'(c_1'c_2')^{\deg(a)-1} c_1' \\ &= (c_2'c_1')^{\deg(a)} \end{aligned}$$

But this yields an expression of b as a power of $\sqrt{c_2'c_1'}$, and by Theorem 2.14 such representations are unique; so $\sqrt{b} = \sqrt{c_2'c_1'}$. But then we have

$$|\sqrt{c_2'c_1'}| = |\sqrt{b}| \geq |\sqrt{a}| = |c_1'c_2'| = |c_2'c_1'|$$

Hence $c_2'c_1' = \sqrt{c_2'c_1'} = \sqrt{b}$, as desired. □ Lemma 4.5

An immediate corollary:

Corollary 4.6. *Suppose $a, b \in \Sigma^+$ are conjugate. Then $|\sqrt{a}| = |\sqrt{b}|$ and $\deg(a) = \deg(b)$. In particular, we get that a is primitive if and only if b is.*

We now turn to the promised variation on the first theorem of Lyndon-Schützenberger:

Proposition 4.7. *Suppose $a, b, c \in \Sigma^+$. Then $ab = bc$ if and only if $\deg(a) = \deg(c)$ and there are $d_1, d_2 \in \Sigma^*$ and $n < \omega$ such that*

1. $\sqrt{a} = d_1 d_2$
2. $\sqrt{c} = d_2 d_1$
3. $b = (\sqrt{a})^n d_1$

Again, since $\sqrt{a}$ and $\sqrt{c}$ are primitive, Corollary 2.20 yields that d_1 and d_2 are unique, if they exist.

Proof.

($\Longleftarrow$) Suppose we have such d_1, d_2, and n. Then

$$\begin{aligned} ab &= (\sqrt{a})^{\deg(a)}(\sqrt{a})^n d_1 \\ &= (d_1 d_2)^{\deg(a)}(d_1 d_2)^n d_1 \\ &= (d_1 d_2)^{\deg(a)+n} d_1 \\ &= d_1 (d_2 d_1)^{\deg(a)+n} \\ &= d_1 (d_2 d_1)^n (d_2 d_1)^{\deg(a)} \\ &= (d_1 d_2)^n d_1 (\sqrt{c})^{\deg(c)} \\ &= (\sqrt{a})^n d_1 c \\ &= bc \end{aligned}$$

as desired.

($\Longrightarrow$) Suppose $ab = bc$. Then by the first theorem of Lyndon-Schützenberger (Theorem 2.7) we get that there are $e_1, e_2 \in \Sigma^*$ and $n < \omega$ such that

1. $a = e_1 e_2$
2. $c = e_2 e_1$
3. $b = a^n e_1$

Then by Lemma 4.5 we get that there are $d_1, d_2 \in \Sigma^*$ and $m < \omega$ such that

1. $\sqrt{a} = d_1 d_2$
2. $\sqrt{c} = d_2 d_1$
3. $e_1 = (\sqrt{a})^m d_1$ and in particular $b = a^n e_1 = (\sqrt{a})^{n \deg(a)} (\sqrt{a})^m d_1 = (\sqrt{a})^{n \deg(a) + m} d_1$

That $\deg(a) = \deg(c)$ follows from Corollary 4.6. ☐ Proposition 4.7

As an immediate corollary, we find:

Corollary 4.8. *Suppose $a, b \in \Sigma^*$; let $\varphi(x)$ be $ax = xb$. Then $\varphi(\Sigma^*)$ is either empty or takes the form $(\sqrt{a})^* a_0$ for some proper prefix a_0 of $\sqrt{a}$.*

Our next step is to examine the atomically definable subsets of Σ^*. We begin with a technical lemma:

Lemma 4.9. *Suppose $|a^* a_0 \cap b^* b_0| \geq 2$ with a, b primitive, a_0 a proper prefix of a, and b_0 a proper prefix of b. Then $a = b$ and $a_0 = b_0$.*

Proof. Suppose $|a^*a_0 \cap b^*b_0| \geq 2$. Pick $n_1 < n_2 < \omega$ and $m_1 < m_2 < \omega$ such that $a^{n_1}a_0 = b^{m_1}b_0$ and $a^{n_2}a_0 = b^{m_2}b_0$. But now

$$a^{n_2-n_1}(b^{m_1}b_0) = a^{n_2-n_1}(a^{n_1}a_0) = a^{n_2}a_0 = b^{m_2}b_0 = b^{m_2-m_1}(b^{m_1}b_0)$$

and thus $a^{n_2-n_1} = b^{m_2-m_1}$. But by Theorem 2.14 we have that representations of words as powers of primitives are unique; hence $a = b$. But now a_0 and b_0 are both proper prefixes of $a = b$, and

$$|a_0| = |a^{n_1}a_0| \bmod |a| = |b^{m_1}b_0| \bmod |b| = |b_0|$$

Hence $a_0 = b_0$. □ Lemma 4.9

We now provide a necessary condition for a subset of Σ^* to be definable by an atomic formula with parameters:

Theorem 4.10. *Suppose $A \subseteq \Sigma^*$ is definable by an atomic formula with parameters. Then A is either finite, all of Σ^*, or has finite symmetric difference from e^*e_0 for some primitive $e \in \Sigma^*$ and some proper prefix e_0 of e.*

Proof. Let the atomic formula defining A be

$$a_0xa_1x\cdots xa_n = b_0xb_1x\cdots xb_m$$

where $a_0, \ldots, a_n, b_0, \ldots b_m \in \Sigma^*$. Observe that if $c \in A$ then $|a_0\cdots a_n| + n|c| = |b_0\cdots b_m| + m|c|$; hence if $n \neq m$, then $|c|$ is determined by the above equation, and A is finite. Suppose then that $m = n$; we now apply induction on n.

The base case (i.e. $n = 0$) is immediate.

For the induction step, suppose $0 < n < \omega$. If $A \neq \emptyset$ then one of a_0 and b_0 is a prefix of the other; we may thus reduce the equation to be one of the form

$$a_0xa_1x\cdots xa_n = xb_1x\cdots xb_n$$

i.e. one in which $b_0 = \varepsilon$.

Now, if $|c| \geq |a_0|$ and $c \in A$, then a_0 is a prefix of c; so $c = a_0c'$ where c' is a solution to

$$a_0(a_0x')a_1(a_0x')\cdots(a_0x')a_n = (a_0x')b_1(a_0x')\cdots(a_0x')b_n$$

which reduces to

$$a_0x'a_1a_0x'\cdots a_{n-1}a_0x'a_n = x'b_1a_0x'\cdots b_{n-1}a_0x'b_n$$

But $|b_1a_0| \geq |a_0|$; so we may write $b_1a_0 = d_1d_2$ where $|d_1| = |a_0|$. Then the above equation reduces to the conjunction of

$$a_0x' = x'd_1$$
$$a_1a_0x'a_2a_0x'\cdots a_{n-1}a_0x'a_n = d_2x'b_2a_0x'\cdots b_{n-1}a_0x'b_n$$

Furthermore, as long as $|c| \geq |a_0|$, the converse holds. Note, however, that $|c| \geq |a_0|$ for all but finitely many $c \in \Sigma^*$; hence if B and C are the solution sets to the above two equations, respectively, then A has finite symmetric difference from $a_0(B \cap C)$.

Now, by the induction hypothesis we get that C has the desired form, and by Corollary 4.8 we get that B takes the form $(\sqrt{a_0})^*a_{00}$ for some proper prefix a_{00} of $\sqrt{a_0}$. Now, if C is finite, then A is finite, and we're done; if $C = \Sigma^*$ then A has finite symmetric difference from $a_0(\sqrt{a_0})^*a_{00}$, which is itself a cofinite subset of $(\sqrt{a_0})^*a_{00}$, and we are again done.

Suppose then that C has finite symmetric difference from e^*e_0 for some primitive e and some proper prefix e_0 of e. By Lemma 4.9, if $e \neq \sqrt{a_0}$ or $e_0 \neq a_{00}$ then $B \cap C$ is finite; so A is finite, and we're done. On the other hand, if $e = \sqrt{a_0}$ and $e_0 = a_{00}$ then B and C have finite symmetric difference, and $B \cap C$ has finite symmetric difference from B; so A has finite symmetric difference from $a_0B = a_0(\sqrt{a_0})^*a_{00}$ which is a cofinite subset of $(\sqrt{a_0})^*a_{00}$.

So A has the desired form. □ Theorem 4.10

Remark 4.11. This is by no means a characterization of the atomically definable subsets of Σ^*. To see this, note that if $A \subseteq \Sigma^*$ is defined by

$$a_0 x a_1 x \cdots x a_n = x b_1 x \cdots x b_m$$

then any $c \in A$ is either a prefix of a_0 or has $|c| \geq |a_0|$, which as noted in the above proof implies that $c \in (\sqrt{a_0})^* a_{00}$ for some prefix a_{00} of a_0; in particular, in either case c is a prefix of a_0^n for all sufficiently large n. In further particular, we get that A is a chain under the "is a prefix of" partial order. Hence any finite set that is not a chain cannot be atomically definable.

Even if we restrict to finite sets that are chains, it's not clear which ones are atomically definable. Certainly all singletons are, and there are examples of atomically definable sets with two elements: in $\Sigma = \{0, 1\}$, the solutions to $1101x0x = xx01101$ are $x = 1$ and $x = 1101$.

Happily, for the purposes of determining which sets are quantifier-free definable, these issues vanish, and we can give a true characterization.

Corollary 4.12. *The quantifier-free definable subsets of Σ^* are exactly those that have finite symmetric difference from a set of one of the following forms:*

1. $a_1^* b_1 \cup \cdots \cup a_n^* b_n$ *for primitive $a_1, \ldots, a_n$ and b_i a proper prefix of a_i, or*
2. $(a_1^* b_1 \cup \cdots \cup a_n^* b_n)^c$ *for primitive $a_1, \ldots, a_n$ and b_i a proper prefix of a_i.*

Proof. From Proposition 4.7 we get that all such sets are indeed quantifier-free definable: given $a \in \Sigma^*$ primitive and b a proper prefix of a, if we let $a = bc$ then the equation $ax = xcb$ defines a^*b. From Theorem 4.10, we get that all of the atomically definable subsets of Σ^* take the desired form. It then suffices to check that sets of the above form are closed under Boolean combinations. Closure under complementation is immediate; we check closure under intersection.

Suppose A and B take the form above. There are three cases to consider:

Case 1. If both A and B take the second form, it is clear that their intersection again takes the second form.

Case 2. If A takes the first form and B takes the second form, then their intersection has finite symmetric difference from a set of the form

$$(a_1^* b_1 \cup \cdots \cup a_n^* b_n) \cap (c_1^* d_1)^c \cap \cdots \cap (c_m^* d_m)^c = \bigcup_{i=1}^{n} \left(a_i^* b_i \cap (c_1^* d_1)^c \cap \cdots \cap (c_m^* d_m)^c \right)$$

But by Lemma 4.9 we get either that $a_i^* b_i \cap c_j^* d_j$ is finite or that $a_i = c_j$ and $b_i = d_j$; hence

$$a_i^* b_i \cap (c_1^* d_1)^c \cap \cdots \cap (c_m^* d_m)^c$$

has finite symmetric difference from either $\emptyset$ or $a_i^* b_i$. So $A \cap B$ takes the desired form.

Case 3. If A and B both take the first form, then their intersection has finite symmetric difference from a set of the form

$$\bigcup_{i=1}^{n} \bigcup_{j=1}^{m} (a_i^* b_i \cap c_j^* d_j)$$

Again by Lemma 4.9 we get that each $a_i^* b_i \cap c_j^* d_j$ is either finite or $a_i^* b_i$; hence $A \cap B$ takes the desired form. ☐ Corollary 4.12

Remark 4.13. As noted in Theorem 2.7, if a is primitive and b is a prefix of a, say with $a = bc$, and $n < \omega$, then $a^n b = (bc)^n b = b(cb)^n$; furthermore, since a is primitive, Corollary 4.6 yields that cb is primitive. In particular, this means that $a^* b = b(cb)^*$, where cb is primitive and b is a suffix of cb. We thus see that the choice to represent the quantifier-free definable subsets of Σ^* using $a^* b$ for a primitive and b a prefix of a was somewhat arbitrary; we could just as easily have represented them using ba^* for a primitive and b a suffix of a.

4.2 Quantifier-free definable subsets of Σ^* and star-free languages

A natural question at this point is whether the class of quantifier-free definable subsets of Σ^* fits into the formal languages hierarchy, and if so where. As it turns out, it does: it falls between the class of finite languages and the class of star-free languages. In this section, we will define star-free languages and examine their relation to the quantifier-free definable subsets of Σ^*. A good external reference for star-free languages is [4, Chapter X].

Definition 4.14. The class of *star-free languages* is the smallest class of languages that contains the finite languages and is closed under Boolean combinations and concatenation.

Example 4.15. $0^*1^* \subseteq \{0,1\}^*$ is star-free; to see this, one notes that $0^*1^* = (\emptyset^c 10 \emptyset^c)^c$. More generally, the set of strings that don't contain a fixed word as a substring is star-free.

Remark 4.16. It follows from the definition and Proposition 2.28 that the class of regular languages contains the finite languages and is closed under Boolean combinations and concatenation; hence every star-free language is regular.

Proposition 4.17. *Suppose $A \subseteq \Sigma^*$ has finite symmetric difference from $B \subseteq \Sigma^*$ and B is star-free. Then A is star-free.*

Proof. Write $A = (B \setminus (B \setminus A)) \cup (A \setminus B)$. But B is star-free, and $B \setminus A$ and $A \setminus B$ are finite, and thus star-free; so A is a Boolean combination of star-free languages, and is thus star-free. ☐ Proposition 4.17

We now turn to the relation between the quantifier-free definable subsets of Σ^* and the star-free languages.

Proposition 4.18. *Suppose $a \in \Sigma^*$ is primitive. Then a^* is star-free.*

We will first need a lemma.

Lemma 4.19. *Suppose $a \in \Sigma^+$ is primitive; suppose $a^n = b_1 a b_2$ for some $n < \omega$ and $b_1, b_2 \in \Sigma^*$. Then b_1 and b_2 are powers of a.*

Proof. Let $c_1 = b_1$ and $c_2 = ab_2$. Then $a^n = c_1c_2$ and $ab_2b_1 = c_2c_1$; so, by Corollary 4.6, we get that $|\sqrt{ab_2b_1}| = |\sqrt{a^n}| = |a|$. So $\sqrt{ab_2b_1}$ is the prefix of ab_2b_1 of length $|a|$, which is just a; so $\sqrt{ab_2b_1} = a$, and $ab_2b_1 = a^n$. Now by Lemma 4.5, we find that if $c_1 = a^k c_1'$ with $|c_1'| < |a|$, then there is c_2' such that $c_1'c_2' = a = \sqrt{ab_2b_1} = c_2'c_1'$. Since a is primitive, Proposition 2.19 yields that $c_1' = \varepsilon$, and thus that $b_1 = a^k$ is a power of a. But then $a^n = b_1ab_2 = a^{k+1}b_2$; so $b_2 = a^{n-k-1}$ is also a power of a. ☐ Lemma 4.19

Proof of Proposition 4.18. Let

$$B = \{ b \in \Sigma^* : |b| \leq |a|, b \text{ is not a prefix of } a \}$$
$$C = (a\Sigma^*) \cap (\Sigma^* a) \cap \left(\Sigma^* a B \Sigma^*\right)^c$$

In words:

- B is the set of strings of length at most $|a|$ that are not prefixes of a.
- C is the set of strings $c \in \Sigma^*$ satisfying the following:
  1. c starts with a.
  2. c ends with a.
  3. Whenever a occurs as a substring of c, it is not followed by an element of B; i.e. it is either followed by a, or it is followed by some prefix of a after which the string (that is, c) ends.

Now, B is finite, and thus star-free; we then get that C is star-free since it is given as a Boolean combination and concatenation of star-free languages. (Note that $\Sigma^* = \emptyset^c$ is star-free.)

Claim 4.20. $C = a^+$.

Proof.

($\subseteq$) Suppose $c \in C$; we will show by induction on $|c|$ that $c \in a^+$.

For the base case, suppose $|c| < 2|a|$. Since $c \in C$, we get that c begins and ends with a; say $c = ad_1 = d_2a$. It is then clear that $|d_1| = |d_2| < |a|$. Now, by the third condition for membership in C, we find that d_1 is a prefix of a; say $a = d_1e$. Substituting, we see $c = d_1ed_1 = d_2d_1e$; thus, since $|d_1| = |d_2|$, we have $d_1 = d_2$. So $c = ad_1 = d_1a$. But a is primitive; so, by Remark 2.13, we get that $d_1 \in a^*$. Since $|d_1| < |a|$, we then get that $d_1 = \varepsilon$, and $c = a \in a^+$.

For the induction step, suppose $|c| \geq 2|a|$. Then by the first and third conditions on membership in C, we find that c begins with aa; say $c = aad$. I now claim that $ad \in C$; to see this, we verify the conditions:

1. ad clearly begins with a.
2. ad ends with a since $c = aad$ ends with a and ad is a suffix of c of length $\geq |a|$.
3. Any occurrence of a as a substring of ad is also an occurrence of a as a substring of c; furthermore, since ad is a suffix of c, the strings following said occurrence must also follow the corresponding occurrence in c, and thus cannot lie in B.

So $ad \in C$. But $|ad| < |aad| = c$; so, by the inductive hypothesis, we get that $ad \in a^+$. Hence $c = aad \in a^+$ as well.

($\supseteq$) Suppose we are given $c = a^n$ for $0 < n < \omega$. It is clear that c starts and ends with a; it remains to check the third condition.

Suppose then that $c = a^n = b_1ab_2$. Then by Lemma 4.19 we get that b_2 is a power of a. In particular, every string of length $\leq |a|$ that follows the occurrence of a in question must be a prefix of a; so the occurrence of a in question is not followed by an element of B.

So c satisfies the third condition. So $c \in C$. ☐ Claim 4.20

So a^+ is star-free; so $a^* = \varepsilon \cup a^+$ is also star-free. ☐ Proposition 4.18

It is worth remarking that Lemma 4.19 fails if a is not primitive; consider for example $a = 00$, $b_1 = 0$, and $b_2 = 0$. Indeed, Proposition 4.18 fails if a is not primitive; as noted in [2, Section 3], we have that $(\ell\ell)^*$ is not star-free for any $\ell \in \Sigma$.

Corollary 4.21. *Every quantifier-free definable subset of Σ^* is star-free, and in particular is regular.*

Proof. By Proposition 4.18 we get that a^* is star-free if $a \in \Sigma^*$ is primitive; if a_0 is a proper prefix of a, we then get that a^*a_0 is star-free as the concatenation of star-free languages. It then follows that Boolean combinations of sets of the form a^*a_0 are star-free, where a is primitive and a_0 is a proper prefix of a. By Proposition 4.17, we then have that any set having finite symmetric difference from a Boolean combination of sets of the form a^*a_0 is star-free. But by Corollary 4.12 all quantifier-free definable subsets of Σ^* take this form; so all quantifier-free definable subsets of Σ^* are star-free. ☐ Corollary 4.21

Remark 4.22. This provides another proof that $\mathrm{Th}(\Sigma^*, \varepsilon, \cdot)$ does not admit quantifier elimination: as noted in Example 4.2, we have that $\{ a^2 : a \in \Sigma^* \}$ is definable but not regular, and hence not star-free, and hence not quantifier-free definable.

One might ask whether the converse holds, and the quantifier-free definable subsets of Σ^* are exactly the star-free languages; it does not.

Example 4.23. Let $A = 0^*1^* \subseteq \{0, 1\}^*$. As pointed out in Example 4.15, we have that A is star-free; I claim it is not quantifier-free definable. To see this, we will use Corollary 4.12: we check that it does not have finite symmetric difference from a set of either of the specified forms.

1. We first consider sets of the form $B = a_1^*b_1 \cup \cdots \cup a_n^*b_n$ for primitive $a_1, \ldots, a_n$ and b_i a proper prefix of a_i.

   Case 1. Suppose one of the a_i contains both a 0 and a 1. Then $a_i^*b_i$ contains infinitely many strings with an occurrence of 10, none of which lie in A; hence A and B do not have finite symmetric difference.

Case 2. Suppose none of the a_i contains both a 0 and a 1. Then each a_i is primitive and is composed either entirely of 0 or entirely of 1; hence each a_i is either 0 or 1, and each $a_i^* b_i$ is either 0^* or 1^*. In particular, every word of the form $0^n 1^n$ for $0 < n < \omega$ is in $A \setminus B$; so A and B do not have finite symmetric difference.

2. We now consider sets of the form $B = (a_1^* b_1)^c \cap \cdots \cap (a_n^* b_n)^c$ for primitive $a_1, \ldots, a_n$ and b_i a proper prefix of a_i. Note that whenever $0 < m < \omega$ we have that $1^m 0^m \notin A$; however, one can check that if $1^m 0^m \in a^* b$ for some $a \in \Sigma^*$ and some prefix b of a, then one of a or b is equal to $1^m 0^m$. Hence, with the exception of the finitely many m with $1^m 0^m$ equal to one of the a_i or b_i, we have that $1^m 0^m \in B$. Hence $B \setminus A$ is infinite, and A and B do not have finite symmetric difference. □

We have thus seen how the quantifier-free definable subsets of Σ^* fit into the formal languages hierarchy: they properly contain the finite languages and are properly contained in the star-free languages (which are themselves properly contained in the regular languages).

Remark 4.24. The class of definable subsets of Σ^* contains the finite languages and is closed under Boolean combinations; furthermore, as noted in Proposition 4.1, it is closed under concatenation. Hence it contains the class of star-free languages; i.e. every star-free language is definable.

5 A model-theoretic interpretation of regularity

In this section, we present a model-theoretic interpretation of regularity.

5.1 φ-types and the rank-zero case

We begin by recalling the local theory of types and rank; we will later interpret regularity using this local theory. We mostly follow the treatment sketched in [3, Chapter 1], though said treatment sometimes assumes stability, which fails here. (Indeed, the φ we will use will not in general be stable; see Example 5.35.)

Throughout this subsection, we fix a language L, a complete theory T, and an L-formula $\varphi(\overline{x}; \overline{u})$. We let $\mathfrak{C}$ denote a fixed sufficiently saturated model of T.

Definition 5.1. Suppose A is a parameter set.

- A *φ-instance* is a formula of the form $\varphi(\overline{x}; \overline{a})$ for some tuple $\overline{a}$ from $\mathfrak{C}$.
- A *φ-formula over A* is a formula $\psi(\overline{x})$ with parameters from A that is equivalent to some Boolean combination of φ-instances. (We do not require that the φ-instances have parameters coming from A.)
- A *complete φ-type over A* is a maximally consistent set of φ-formulae over A. We use $S_\varphi(A)$ to denote the set of complete φ-types over A.
- Suppose $\overline{a}$ is a tuple from $\mathfrak{C}$ of the same arity as $\overline{x}$. We define the *φ-type of $\overline{a}$ over A* to be

$$\mathrm{tp}_\varphi(\overline{a}/A) = \{\, \psi(\overline{x}) : \psi \text{ is a } \varphi\text{-formula over } A, \mathfrak{C} \models \psi(\overline{a}) \,\}$$

- Suppose $B \subseteq \mathfrak{C}$. (We do not require that B be small.) We say a complete φ-type over A is *realized* in B if it takes the form $\mathrm{tp}_\varphi(\overline{a}/A)$ for some tuple $\overline{a}$ from B.

Remark 5.2. We can endow $S_\varphi(A)$ with a topology in the usual way: we declare the basic open sets to be those of the form $[\psi] = \{\, p \in S_\varphi(A) : \psi \in p \,\}$ where $\psi(\overline{x})$ is a φ-formula. This topology is compact and Hausdorff, and the basic open sets are also closed.

Remark 5.3. If $A \subseteq B$ then any complete φ-type over A extends to at least one complete φ-type over B; furthermore, different φ-types over A cannot extend to the same φ-type over B. Hence $|S_\varphi(A)| \leq |S_\varphi(B)|$.

The following proposition tells us that when working over a model, the notion of φ-types reduces to sets of φ-instances.

Proposition 5.4. *Suppose $\mathcal{M} \models T$; suppose $p, q \in S_\varphi(M)$. Suppose that for all tuples $\overline{c}$ from M we have*

$$\varphi(\overline{x}; \overline{c}) \in p \iff \varphi(\overline{x}; \overline{c}) \in q$$

Then $p = q$.

Proof. Suppose $\psi(\overline{x})$ is a φ-formula over M; say it is equivalent to some Boolean combination $\psi'(\overline{x}; \overline{d})$ of φ-instances. Then

$$\mathfrak{C} \models \exists \overline{u} \forall \overline{x} (\psi(\overline{x}) \leftrightarrow \psi'(\overline{x}; \overline{u}))$$

But $\mathfrak{C} \succeq \mathcal{M}$; so $\mathcal{M}$ does as well. So there is a tuple $\overline{d'}$ from M such that $\psi(\overline{x})$ is equivalent to $\psi'(\overline{x}; \overline{d'})$. By hypothesis, we get $\psi'(\overline{x}; \overline{d'}) \in p$ if and only if $\psi'(\overline{x}; \overline{d'}) \in q$. Since p and q are maximally consistent, we then get that $\psi(\overline{x}) \in p$ if and only if $\psi(\overline{x}) \in q$. So $p = q$. □ Proposition 5.4

Note in particular that Proposition 5.4 holds when $\mathcal{M} = \mathfrak{C}$.

We now give an exposition of Cantor-Bendixson rank, which we will eventually use to define φ-rank.

Definition 5.5. Suppose X is a topological space; suppose α is an ordinal. We define the α^{th} *Cantor-Bendixson derivative* of X, denoted $X^{(\alpha)}$, via transfinite recursion:

- $X^{(0)} = X$.
- Given a successor ordinal $\alpha + 1$, we set $X^{(\alpha+1)} = \{\, p \in X^{(\alpha)} : p \text{ is not isolated in } X^{(\alpha)} \,\}$.
- Given a limit ordinal β, we set
$$X^{(\beta)} = \bigcap_{\alpha < \beta} X^{(\alpha)}$$

Remark 5.6. One can check by transfinite induction that the following hold:

1. $X^{(\alpha)}$ is a closed subset of X for all α.
2. If $X \subseteq Y$ then $X^{(\alpha)} \subseteq Y^{(\beta)}$.
3. $X^{(\alpha)} \supseteq X^{(\beta)}$ if $\alpha \leq \beta$.
4. If $U \subseteq X$ is open then $U^{(\alpha)} = X^{(\alpha)} \cap U$.

Proposition 5.7. *Suppose X is compact; suppose*

$$\bigcap_\alpha X^{(\alpha)} = \emptyset$$

Then there is a largest α with $X^{(\alpha)} \neq \emptyset$.

Proof. By Remark 5.6, the $X^{(\alpha)}$ are all closed in X; so, by compactness, we get that the collection of $X^{(\alpha)}$ must not satisfy the finite intersection property, and there is some β such that $X^{(\beta)} = \emptyset$. Without loss of generality, assume β is minimal. Now, if β were a limit ordinal, then we would have

$$\emptyset = X^{(\beta)} = \bigcap_{\alpha < \beta} X^{(\alpha)}$$

and again by compactness we would have $X^{(\alpha)} = \emptyset$ for some $\alpha < \beta$, a contradiction. So $\beta = \alpha + 1$ for some α, and α is our desired maximum. □ Proposition 5.7

Definition 5.8. Given a compact space X, we define the *Cantor-Bendixson rank of X*, denoted $\mathrm{CB}(X)$, to be ∞ if

$$\bigcap_\alpha X^{(\alpha)} \neq \emptyset$$

and otherwise to be the largest α with $X^{(\alpha)} \neq \emptyset$.

Remark 5.9. Suppose X is compact and has $\mathrm{CB}(X) < \infty$. It follows from Remark 5.6 that if α is an ordinal then $X^{(\alpha)} \neq \emptyset$ if and only if $\alpha \leq \mathrm{CB}(X)$.

Remark 5.10. Suppose X is compact and has $\mathrm{CB}(X) < \infty$. Since $X^{(\mathrm{CB}(X)+1)} = \emptyset$, we get that every point in $X^{(\mathrm{CB}(X))}$ is isolated. But Remark 5.6 yields that $X^{(\mathrm{CB}(X))}$ is a closed subset of X, and is thus compact. But now $\{\, \{p\} : p \in X^{(\mathrm{CB}(X))} \,\}$ is an open cover of $X^{(\mathrm{CB}(X))}$ with no proper subcover; by compactness, it is finite. So $X^{(\mathrm{CB}(X))}$ is finite.

Definition 5.11. Suppose X is a compact space with $\mathrm{CB}(X) < \infty$. We define the *Cantor-Bendixson multiplicity of* X, denoted $\text{CB-mult}(X)$, to be $|X^{(\mathrm{CB}(X))}|$.

We are now in a position to define local ranks:

Definition 5.12. Suppose $\psi(\overline{x})$ is a φ-formula. We define the *φ-rank of ψ* to be $R_\varphi(\psi) = \mathrm{CB}([\psi]) = \mathrm{CB}(\{\, p \in S_\varphi(\mathfrak{C}) : \psi \in p \,\})$. (Note that since $[\psi]$ is a basic open set of $S_\varphi(\mathfrak{C})$, it is also closed; it is thus compact, since $S_\varphi(\mathfrak{C})$ is compact, and thus has a Cantor-Bendixson rank.) If $R_\varphi(\psi) < \infty$, we define the *φ-multiplicity of ψ* to be $\mathrm{mult}_\varphi(\psi) = \text{CB-mult}([\psi])$.

It's not immediately apparent that $R_\varphi(\psi)$ encodes anything meaningful from a logical standpoint. The following proposition phrases R_φ in terms of branching, whence it is clearer that R_φ is a local analogue of Morley rank:

Proposition 5.13. *Suppose $\psi(\overline{x})$ is a φ-formula and α is an ordinal. Then $R_\varphi(\psi) \geq \alpha + 1$ if and only if there are φ-formulae $(\psi_i(\overline{x}) : i < \omega)$ satisfying:*

1. *The ψ_i are pairwise inconsistent.*
2. *$\mathfrak{C} \models \forall \overline{x}(\psi_i(\overline{x}) \to \psi(\overline{x}))$ for all $i < \omega$.*
3. *$R_\varphi(\psi_i) \geq \alpha$ for all $i < \omega$.*

Proof.

($\implies$) Suppose $R_\varphi(\psi) = \mathrm{CB}([\psi]) \geq \alpha + 1$; then there is $p(\overline{x}) \in [\psi]^{(\alpha+1)}$. In particular, p is not isolated in $[\psi]^{(\alpha)}$. We inductively define φ-formulae $\xi_i(\overline{x})$ such that if

$$\chi_i(\overline{x}) = \bigwedge_{j<i} \neg\xi_i(\overline{x})$$

then for all $i < \omega$ we have the following:

1. $R_\varphi(\psi \wedge \xi_i \wedge \chi_i) \geq \alpha$.
2. $\chi_{i+1} \in p$.

Suppose we have defined ξ_j for $j < i$. Since $\chi_i \in p$, we get that $[\psi]^{(\alpha)} \cap [\chi_i]$ is an open neighbourhood of p in $[\psi]^{(\alpha)}$; since p is not isolated in $[\psi]^{(\alpha)}$, there is some $q(\overline{x}) \in [\psi]^{(\alpha)} \cap [\chi_i]$ such that $q \neq p$. Since $q \neq p$, there is some φ-formula $\xi_i(\overline{x})$ with $\xi_i \in q$ and $\neg\xi_i \in p$. We verify that ξ_i satisfies the desired properties:

1. Note that by Remark 5.6 we have $q \in [\psi]^{(\alpha)} \cap [\psi \wedge \xi_i \wedge \chi_i] = [\psi \wedge \xi_i \wedge \chi_i]^{(\alpha)}$; so $R_\varphi(\psi \wedge \xi_i \wedge \chi_i) = \mathrm{CB}([\psi \wedge \xi_i \wedge \chi_i]) \geq \alpha$.
2. By choice of ξ_i we have $\neg\xi_i \in p$; by the induction hypothesis we have that $\chi_i \in p$. Hence $\chi_{i+1} = \chi_i \wedge \neg\xi_i \in p$.

(Note that the above construction works perfectly well when $i = 0$.) For $i < \omega$, we now let $\psi_i(\overline{x}) = \psi(\overline{x}) \wedge \xi_i(\overline{x}) \wedge \chi_i(\overline{x})$. We verify that the ψ_i satisfy the desired properties:

1. Note that if $i < j < \omega$ then $\chi_j \wedge \xi_i$ is inconsistent; hence $\psi_i \wedge \psi_j$ is inconsistent.
2. It is clear that $\mathfrak{C} \models \forall \overline{x}(\psi_i(\overline{x}) \to \psi(\overline{x}))$ for $i < \omega$.

3. By construction we have that $R_\varphi(\psi_i) \geq \alpha$.

So we have our desired $(\psi_i : i < \omega)$.

$(\Longleftarrow)$ Suppose we have such $(\psi_i : i < \omega)$. Suppose $i < \omega$. Since $\mathrm{CB}([\psi_i]) = R_\varphi(\psi_i) \geq \alpha$, there is some $p_i(\overline{x}) \in [\psi_i]^{(\alpha)}$. Since $\mathfrak{C} \models \forall \overline{x}(\psi_i(\overline{x}) \to \psi(\overline{x}))$, Remark 5.6 yields that $[\psi_i]^{(\alpha)} \subseteq [\psi]^{(\alpha)}$; so $p_i \in [\psi]^{(\alpha)}$.

Now, since the ψ_i are pairwise inconsistent and $\psi_i \in p_i$, we get that the p_i are distinct. So $[\psi]^{(\alpha)}$ is infinite. But $[\psi]$ is compact, and Remark 5.6 tells us that $[\psi]^{(\alpha)}$ is a closed subset of $[\psi]$, and is thus compact. So every infinite subset of $[\psi]^{(\alpha)}$ has an accumulation point, and in particular $\{ p_i : i < \omega \}$ has an accumulation point; say $p \in [\psi]^{(\alpha)}$. But then p is not isolated in $[\psi]^{(\alpha)}$; so $p \in [\psi]^{(\alpha+1)}$. So $R_\varphi(\psi) = \mathrm{CB}([\psi]) \geq \alpha + 1$. □ Proposition 5.13

We have a similar characterization of $\mathrm{mult}_\varphi(\psi)$ that is reminiscent of Morley degree:

Proposition 5.14. *Suppose $\psi(\overline{x})$ is a φ-formula with $R_\varphi(\psi) < \infty$. Then $\mathrm{mult}_\varphi(\psi)$ is the largest $n < \omega$ such that there are $(\psi_i : i < n)$ satisfying:*

1. *The ψ_i are pairwise inconsistent.*
2. *$\mathfrak{C} \models \forall \overline{x}(\psi_i(\overline{x}) \to \psi(\overline{x}))$ for all $i < n$.*
3. *$R_\varphi(\psi_i) \geq R_\varphi(\psi)$ for all $i < n$.*

Proof. Let $\alpha = R_\varphi(\psi)$.

We first construct $(\psi_p : p \in [\psi]^{(\alpha)})$ satisfying the desired properties. (Recall that $\mathrm{mult}_\varphi(\psi) = \text{CB-mult}([\psi]) = \left|[\psi]^{(\alpha)}\right|$.) By definition we have $[\psi]^{(\alpha+1)} = \emptyset$, and hence that every $p \in [\psi]^{(\alpha)}$ is isolated in $[\psi]^{(\alpha)}$. For $p \in [\psi]^{(\alpha)}$, fix a φ-formula $\chi_p(\overline{x})$ that isolates p in $[\psi]^{(\alpha)}$; let

$$\psi_p(\overline{x}) = \psi(\overline{x}) \wedge \chi_p(\overline{x}) \wedge \bigwedge_{\substack{q \in [\psi]^{(\alpha)} \\ q \neq p}} \neg \chi_q(\overline{x})$$

It is clear that the ψ_p are pairwise inconsistent and that $\mathfrak{C} \models \forall \overline{x}(\psi_p(\overline{x}) \to \psi(\overline{x}))$ for all $p \in [\psi]^{(\alpha)}$. Finally, if $p \in [\psi]^{(\alpha)}$, then Remark 5.6 yields that $p \in [\psi]^{(\alpha)} \cap [\psi_p] = [\psi_p]^{(\alpha)}$. Hence $R_\varphi(\psi_p) = \mathrm{CB}([\psi_p]) \geq \alpha = R_\varphi(\psi)$. So the $(\psi_p : p \in [\psi]^{(\alpha)})$ satisfy the desired properties.

Suppose now that we have $(\psi_i : i < n)$ satisfying the desired properties; we will show that $n \leq \mathrm{mult}_\varphi(\psi)$. Suppose $i < n$. Since $R_\varphi(\psi_i) \geq \alpha$, there is $p_i(\overline{x}) \in [\psi_i]^{(\alpha)}$; since $\mathfrak{C} \models \forall \overline{x}(\psi_i(\overline{x}) \to \psi(\overline{x}))$, Remark 5.6 yields that $[\psi_i]^{(\alpha)} \subseteq [\psi]^{(\alpha)}$, and hence that $p_i \in [\psi]^{(\alpha)}$. Since the ψ_i are pairwise inconsistent and $\psi_i \in p_i$, we get that the p_i are distinct. So $n \leq \left|[\psi]^{(\alpha)}\right| = \mathrm{mult}_\varphi(\psi)$. □ Proposition 5.14

We are interested in the case where $R_\varphi(\overline{x} = \overline{x}) = 0$. (Note that $\overline{x} = \overline{x}$ is a φ-formula: for any tuple $\overline{a}$ from $\mathfrak{C}$, we have that $\overline{x} = \overline{x}$ is equivalent to $\varphi(\overline{x}; \overline{a}) \vee \neg\varphi(\overline{x}; \overline{a})$.) Recall that globally we have $\mathrm{RM}(\overline{x} = \overline{x}) = 0$ if and only if the domain is finite; hence we are examining some local analogue of having finite domain.

The following notation will be convenient:

Notation 5.15. Suppose $\mathcal{M} \models T$. We let $S^{\mathrm{R}}_\varphi(M) = \{ p \in S_\varphi(M) : p \text{ is realized in } M \}$.

Lemma 5.16. *Suppose $\mathcal{M} \models T$ and $S^{\mathrm{R}}_\varphi(M)$ is finite. Then $S_\varphi(M) = S^{\mathrm{R}}_\varphi(M)$.*

Proof. Suppose for contradiction that we had $p \in S_\varphi(M) \setminus S^{\mathrm{R}}_\varphi(M)$. Then for any $q \in S^{\mathrm{R}}_\varphi(M)$, we have $q \neq p$; so there is a φ-formula $\psi_q(\overline{x})$ in $p \setminus q$. Let

$$\psi(\overline{x}) = \bigwedge_{q \in S^{\mathrm{R}}_\varphi(\mathfrak{C})} \psi_q(\overline{x})$$

(Note that the conjunction is finite since $S^{\mathrm{R}}_\varphi(M)$ is finite.) Then for any $q \in S^{\mathrm{R}}_\varphi(M)$ we have that $q \cup \{ \psi \}$ is inconsistent; so $\psi(\overline{x})$ is not satisfiable in $\mathcal{M}$. But $\psi \in p$; so p is not finitely satisfiable in $\mathcal{M}$, and p is not consistent, a contradiction. So no such p exists, and $S_\varphi(M) = S^{\mathrm{R}}_\varphi(M)$. □ Lemma 5.16

The case $R_\varphi(\overline{x} = \overline{x}) = 0$ is characterized by the following proposition:

Proposition 5.17. *The following are equivalent:*

1. $R_\varphi(\overline{x} = \overline{x}) = 0$.
2. *There is* $\mathcal{M} \models T$ *such that* $S_\varphi(M)$ *is finite.*
3. *There is* $\mathcal{M} \models T$ *such that* $S^{\mathrm{R}}_\varphi(M)$ *is finite.*
4. *There is* $n < \omega$ *such that whenever* $\mathcal{M} \models T$ *we have* $|S_\varphi(M)| = n$.
5. *There is* $n < \omega$ *such that whenever* $\mathcal{M} \models T$ *we have* $|S^{\mathrm{R}}_\varphi(M)| = n$.

Proof.

(1) $\implies$ (2) Suppose $0 = R_\varphi(\overline{x} = \overline{x}) = \mathrm{CB}(S_\varphi(\mathfrak{C}))$. Then by Remark 5.10 we have that $(S_\varphi(\mathfrak{C}))^{(0)} = S_\varphi(\mathfrak{C})$ is finite.

(2) $\implies$ (3) Immediate.

(3) $\implies$ (5) Suppose there is $\mathcal{M} \models T$ with $|S^{\mathrm{R}}_\varphi(M)| = n < \aleph_0$; suppose $\mathcal{N} \models T$. We will show that $|S^{\mathrm{R}}_\varphi(N)| = n$.

Note that given $\mathcal{D} \models T$ and tuples $\overline{a}$ and $\overline{b}$ from the domain D of $\mathcal{D}$, Proposition 5.4 yields that $\mathrm{tp}_\varphi(\overline{a}/D) = \mathrm{tp}_\varphi(\overline{b}/D)$ if and only if $\mathcal{D} \models \chi(\overline{a}, \overline{b})$, where

$$\chi(\overline{x}, \overline{y}) = \forall \overline{u}(\varphi(\overline{x}; \overline{u}) \leftrightarrow \varphi(\overline{y}; \overline{u}))$$

Then since $|S^{\mathrm{R}}_\varphi(M)| = n$, we get that

$$\mathcal{M} \models \exists \overline{x_1} \cdots \exists \overline{x_n} \left(\left(\bigwedge_{i \neq j} \neg\chi(\overline{x_i}, \overline{x_j}) \right) \wedge \left(\forall \overline{y} \left(\bigvee_{i=1}^{n} \chi(\overline{x_i}, \overline{y}) \right) \right) \right)$$

But $\mathcal{N} \equiv \mathcal{M}$; so

$$\mathcal{N} \models \exists \overline{x_1} \cdots \exists \overline{x_n} \left(\left(\bigwedge_{i \neq j} \neg\chi(\overline{x_i}, \overline{x_j}) \right) \wedge \left(\forall \overline{y} \left(\bigvee_{i=1}^{n} \chi(\overline{x_i}, \overline{y}) \right) \right) \right)$$

So $|S^{\mathrm{R}}_\varphi(\mathcal{N})| = n$, as desired.

(5) $\implies$ (4) Suppose $\mathcal{M} \models T$; then by hypothesis we have $|S^{\mathrm{R}}_\varphi(M)| = n < \omega$. Lemma 5.16 then yields that $S_\varphi(M) = S^{\mathrm{R}}_\varphi(M)$, and in particular that $|S_\varphi(M)| = n$.

(4) $\implies$ (1) Since in particular $\mathfrak{C} \models T$, the hypothesis yields that $S_\varphi(\mathfrak{C})$ is finite. Then since $S_\varphi(\mathfrak{C})$ is Hausdorff, we get that every $p \in S_\varphi(\mathfrak{C})$ is isolated. So $R_\varphi(\overline{x} = \overline{x}) = \mathrm{CB}(S_\varphi(\mathfrak{C})) = 0$. $\square$ Proposition 5.17

Corollary 5.18. *Suppose* $R_\varphi(\overline{x} = \overline{x}) = 0$. *Then* φ *is stable.*

Proof. Suppose κ is an infinite cardinal; suppose $|B| \leq \kappa$. Proposition 5.17 yields that $S_\varphi(\mathfrak{C})$ is finite; Remark 5.3 then yields that $|S_\varphi(B)| \leq |S_\varphi(\mathfrak{C})| < \aleph_0 \leq \kappa$. $\square$ Corollary 5.18

Remark 5.19. Suppose $R_\varphi(\overline{x} = \overline{x}) = 0$; suppose $\mathcal{M} \models T$. Then by Proposition 5.17 we have that $|S^{\mathrm{R}}_\varphi(M)|$ is finite; Lemma 5.16 then yields that $S^{\mathrm{R}}_\varphi(M) = S_\varphi(M)$. Hence if $R_\varphi(\overline{x} = \overline{x}) = 0$, we have a strong local analogue of saturation.

Remark 5.20. Suppose $R_\varphi(\overline{x} = \overline{x}) = 0$. Then $\mathrm{CB}(S_\varphi(\mathfrak{C})) = 0$, and

$$\mathrm{mult}_\varphi(\overline{x} = \overline{x}) = \text{CB-mult}(S_\varphi(\mathfrak{C})) = |(S_\varphi(\mathfrak{C}))^{(\mathrm{CB}(S_\varphi(\mathfrak{C})))}| = |S_\varphi(\mathfrak{C})|$$

In particular, if $\mathcal{M} \models T$ then Proposition 5.17 and Remark 5.19 yield that

$$\mathrm{mult}_\varphi(\overline{x} = \overline{x}) = |S_\varphi(\mathfrak{C})| = |S_\varphi(M)| = |S^{\mathrm{R}}_\varphi(M)|$$

5.2 Recognizable sets

Our model-theoretic interpretation of regularity will make sense in a more general setting than subsets of Σ^*. In this subsection, will extend the notion of a regular language to subsets of arbitrary monoids; we will later give a model-theoretic characterization of this extension.

Definition 5.21. Suppose M is a monoid; suppose $A \subseteq M$. We say A is a *recognizable subset of* M if there is some finite monoid F and some homomorphism of monoids $\alpha\colon M \to F$ such that A is a union of fibres of α; i.e. $A = \alpha^{-1}(\alpha(A))$.

Remark 5.22. If $A \subseteq M$ is recognizable, say by $\alpha\colon M \to F$, then $\alpha\colon M \to \alpha(M)$ also witnesses that A is recognizable; hence we may equivalently require that α be surjective.

The notion of a recognizable subset of a monoid is a generalization of the notion of a regular language over Σ^*; to see this, we will need the following definitions and result from automata theory:

Definition 5.23. Suppose M is a monoid; suppose $A \subseteq M$.

- We define the *syntactic congruence of* A, denoted θ_A^{Syn}, to be the equivalence relation on M given by $(a, b) \in \theta_A^{\mathrm{Syn}}$ if and only if for all $c_1, c_2 \in M$ we have

$$c_1 a c_2 \in A \iff c_1 b c_2 \in A$$

  Roughly speaking, we demand that we be able to substitute b for a in an expression without changing membership in A.

- We define the *Myhill-Nerode equivalence of* A, denoted θ_A^{MN}, to be the equivalence relation on M given by $(a, b) \in \theta_A^{\mathrm{MN}}$ if and only if for all $c \in M$ we have

$$ac \in A \iff bc \in A$$

  In other words, we demand that we be unable to distinguish a and b by a right-multiplication followed by querying whether the result is in A.

Remark 5.24. θ_A^{Syn} refines θ_A^{MN}.

Remark 5.25. θ_A^{MN} is *right-invariant*: that is, if $(a, b) \in \theta_A^{\mathrm{MN}}$ and $c \in M$, then $(ac, bc) \in \theta_A^{\mathrm{MN}}$.

The following proposition justifies the use of the term "congruence" for θ_A^{Syn}:

Proposition 5.26. *If M is a monoid and $A \subseteq M$, then θ_A^{Syn} is a* congruence*: that is, whenever $a_1, a_2, b_1, b_2 \in M$ satisfy $(a_1, b_1), (a_2, b_2) \in \theta_A^{\mathrm{Syn}}$, then $(a_1 a_2, b_1 b_2) \in \theta_A^{\mathrm{Syn}}$.*

Proof. Suppose $c_1, c_2 \in M$; suppose $c_1 a_1 a_2 c_2 \in A$. Then, since $(a_1, b_1) \in \theta_A^{\mathrm{Syn}}$, we get that $c_1 b_1 a_2 c_2 \in A$; since $(a_2, b_2) \in \theta_A^{\mathrm{Syn}}$, we get that $c_1 b_1 b_2 c_2 \in A$. By symmetry, we get that $c_1 a_1 a_2 c_2 \in A$ if and only if $c_1 b_1 b_2 c_2 \in A$. So $(a_1 a_2, b_1 b_2) \in \theta_A^{\mathrm{Syn}}$, as desired. □ Proposition 5.26

Remark 5.27. If $(a, b) \in \theta_A^{\mathrm{MN}}$ then $a \in A$ if and only if $b \in A$; likewise with θ_A^{Syn}.

If M is a finitely generated free monoid, then these equivalence relations can be used to test whether A is regular:

Proposition 5.28 ([5, Theorems 1 and 2])**.** *Suppose $A \subseteq \Sigma^*$. Then the following are equivalent:*

1. *A is regular.*
2. *θ_A^{Syn} has finitely many congruence classes.*
3. *θ_A^{MN} has finitely many equivalence classes.*

We are now in a position to see that recognizable subsets of a monoid are a true generalization of regular languages over Σ^*.

Proposition 5.29. *Suppose M is a monoid; suppose $A \subseteq M$. Then the following are equivalent:*

1. *A is recognizable.*
2. *θ_A^{Syn} has finitely many congruence classes.*
3. *θ_A^{MN} has finitely many equivalence classes.*

Proof.

(1) $\implies$ (2) Suppose A is recognizable; that is, suppose we have a finite monoid F and a homomorphism $\alpha \colon M \to F$ such that A is a union of fibres of α.

Claim 5.30. *For $a, b \in M$, if $\alpha(a) = \alpha(b)$, then $(a, b) \in \theta_A^{\mathrm{Syn}}$.*

Proof. Suppose $\alpha(a) = \alpha(b)$; suppose $c_1, c_2 \in M$. Then

$$\begin{aligned} c_1 a c_2 \in A &\iff \alpha^{-1}(\alpha(c_1 a c_2)) \subseteq A \\ &\iff \alpha^{-1}(\alpha(c_1)\alpha(a)\alpha(c_2)) \subseteq A \\ &\iff \alpha^{-1}(\alpha(c_1)\alpha(b)\alpha(c_2)) \subseteq A \\ &\iff \alpha^{-1}(\alpha(c_1 b c_2)) \subseteq A \\ &\iff c_1 b c_2 \in A \end{aligned}$$

So $(a, b) \in \theta_A^{\mathrm{Syn}}$, as desired. □ Claim 5.30

Thus, since F is finite, it follows that θ_A^{Syn} has finitely many congruence classes.

(2) $\implies$ (3) This is immediate from Remark 5.24.

(3) $\implies$ (1) Suppose θ_A^{MN} has finitely many equivalence classes. Let $X = M/\theta_A^{\mathrm{MN}}$; then by hypothesis X is finite. (We are considering X as a set; we cannot in general place a quotient monoid structure on X.) Note that any $a \in M$ defines a map $X \to X$ via $b/\theta_A^{\mathrm{MN}} \mapsto ba/\theta_A^{\mathrm{MN}}$; this is well-defined by Remark 5.25. It is clear that the mapping defined by 1_M is the identity mapping, and that the mapping defined by ab is the composition of the mapping defined by b and the mapping defined by a. We have thus defined a homomorphism of monoids $\alpha \colon M \to (X^X)^{\mathrm{op}}$. Note also that if $a, b \in M$ and $\alpha(a) = \alpha(b)$, then $a/\theta_A^{\mathrm{MN}} = \alpha(a)(1_M/\theta_A^{\mathrm{MN}}) = \alpha(b)(1_M/\theta_A^{\mathrm{MN}}) = b/\theta_A^{\mathrm{MN}}$; hence by Remark 5.27 we get that $a \in A$ if and only if $b \in A$. So A is a union of fibres of α. Furthermore, $(X^X)^{\mathrm{op}}$ is finite since X is finite. So A is recognizable. □ Proposition 5.29

Corollary 5.31. *Suppose $A \subseteq \Sigma^*$. Then A is regular if and only if A is recognizable.*

5.3 A model-theoretic characterization

We now present the promised characterization of recognizability in terms of local types and local rank. To do so, we will need to change our language from L_{Mon} to the following:

Definition 5.32. We work in the language of monoids expanded by a unary predicate; i.e. $L_P = \{1, \cdot, P\}$ where

- 1 is a constant symbol.
- $\cdot$ is a binary function symbol.
- P is a unary predicate symbol.

Remark 5.33. Given a monoid M and $A \subseteq M$, we can view (M, A) as an L_P-structure by interpreting 1 and $\cdot$ as the monoid structure and interpreting P as A.

Throughout the rest of this section, we fix a monoid M and a (not necessarily recognizable) subset A; we work in $\mathrm{Th}(M, A)$.

Definition 5.34. We define L_P-formulae

$$\varphi^{\mathrm{Syn}}(x;u,v) = P(uxv)$$
$$\varphi^{\mathrm{MN}}(x;v) = P(xv)$$

Unfortunately, φ^{Syn} and φ^{MN} aren't in general stable:

Example 5.35. Let $M = \{0,1\}^*$; let $A = \{0^n 1^m : n < m < \omega\} \subseteq M$. For $i < \omega$, let $a_i = 0^i$, $b_i = 1^i$, and $c_i = \varepsilon$. Then for $i, j < \omega$ we have

$$(M,A) \models \varphi^{\mathrm{MN}}(a_i; b_j) \iff 0^i 1^j \in A \iff i < j$$

and

$$(M,A) \models \varphi^{\mathrm{Syn}}(a_i; c_j, b_j) \iff 0^i 1^j \in A \iff i < j$$

So φ^{MN} and φ^{Syn} have the order property, and are not stable.

We can now deduce our long-awaited characterization of recognizability:

Theorem 5.36. *The following are equivalent:*

1. *A is recognizable.*
2. *$R_{\varphi^{\mathrm{Syn}}}(x = x) = 0$.*
3. *$R_{\varphi^{\mathrm{MN}}}(x = x) = 0$.*

Proof. Observe that for $a, b \in M$ we have $(a,b) \in \theta_A^{\mathrm{Syn}}$ if and only if for all $c_1, c_2 \in M$ we have

$$(M,A) \models \varphi^{\mathrm{Syn}}(a; c_1, c_2) \leftrightarrow \varphi^{\mathrm{Syn}}(b; c_1, c_2)$$

By Proposition 5.4, this is equivalent to requiring that $\mathrm{tp}_{\varphi^{\mathrm{Syn}}}(a/M) = \mathrm{tp}_{\varphi^{\mathrm{Syn}}}(b/M)$. In particular, θ_A^{Syn} has finitely many congruence classes if and only if $S^{\mathrm{R}}_{\varphi^{\mathrm{Syn}}}(M)$ is finite; by Proposition 5.17, this is equivalent to requiring that $R_{\varphi^{\mathrm{Syn}}}(x = x) = 0$.

In an identical way, we get that θ_A^{MN} has finitely many equivalence classes if and only if $R_{\varphi^{\mathrm{MN}}}(x = x) = 0$. Proposition 5.29 then yields the desired result. ☐ Theorem 5.36

Corollary 5.37. *If A is recognizable then φ^{Syn} and φ^{MN} are stable.*

Proof. This follows from Theorem 5.36 and Corollary 5.18. ☐ Corollary 5.37

Theorem 5.38. *Suppose A is recognizable. Then $\mathrm{mult}_{\varphi^{\mathrm{Syn}}}(x = x)$ is the number of congruence classes of θ_A^{Syn} and $\mathrm{mult}_{\varphi^{\mathrm{MN}}}(x = x)$ is the number of equivalence classes of θ_A^{MN}.*

Proof. This follows from Remark 5.20 and the proof of Theorem 5.36. ☐ Theorem 5.38

These numbers in fact provide information about how complex an automaton recognizing A must be: without going into any detail, we will simply mention that the former is the size of the syntactic monoid of A, and the latter is the state complexity of A. (See [4, Section IV.4.1] and [7, Sections 3.9 and 3.11], respectively.)

It follows from Theorem 5.36 and Remark 5.19 that if A is recognizable then every complete φ^{Syn}-type over M is realized in M; likewise with φ^{MN}. We end by giving an example showing that the converse is false: we exhibit a non-recognizable set for which every complete φ^{MN}-type over M is realized in M.

Example 5.39. Let $M = \{0,1,2\}^*$. Given $a \in M$ and $i \in \{0,1,2\}$, we let $n_i(a)$ denote the number of occurrences of i in a. Now, let $A = \{a \in \{0,1\}^* : n_0(a) = n_1(a)\} \subseteq M$; let $T = \mathrm{Th}(M, A)$. Then $(0^n : n < \omega)$ are pairwise unrelated by θ_A^{MN}, so Proposition 5.29 yields that A is not recognizable; however, we will see that every complete φ^{MN}-type over M is realized in M.

Claim 5.40. *Suppose $a \in \{0,1\}^* \subseteq M$. Then $\varphi^{\mathrm{MN}}(x; a)$ isolates a φ^{MN}-type over M.*

Proof. For $b \in M$, let $\Delta(b) = n_1(b) - n_0(b)$. Note that for any $b, c \in M$ we have $(M, A) \models \varphi^{\mathrm{MN}}(b; c)$ if and only if $b, c \in \{0,1\}^*$ and $\Delta(b) + \Delta(c) = 0$. In particular, we get that $\varphi^{\mathrm{MN}}(x; a) \wedge \varphi^{\mathrm{MN}}(x; a')$ is consistent if and only if $a' \in \{0,1\}^*$ and $\Delta(a) = \Delta(a')$. Hence Proposition 5.4 yields that $\varphi^{\mathrm{MN}}(x; a)$ isolates a φ^{MN}-type over M; namely, the one determined by

$$\{\varphi^{\mathrm{MN}}(x; a') : a' \in \{0,1\}^*, \Delta(a) = \Delta(a')\} \cup \{\neg\varphi^{\mathrm{MN}}(x; a') : a' \notin \{0,1\}^* \text{ or } \Delta(a) \neq \Delta(a')\}$$

□ Claim 5.40

Hence if $p(x) \in S_{\varphi^{\mathrm{MN}}}(M)$ contains $\varphi^{\mathrm{MN}}(x; a)$ for some $a \in \{0,1\}^*$, then p is isolated, and thus realized in M. Furthermore, if $a \in M \setminus \{0,1\}^*$, then $\varphi^{\mathrm{MN}}(x; a)$ is inconsistent; hence if $p(x) \in S_{\varphi^{\mathrm{MN}}}(M)$ contains no formula of the form $\varphi^{\mathrm{MN}}(x; a)$ for some $a \in \{0,1\}^*$, then p is the complete φ^{MN}-type over M determined by

$$\{\neg\varphi^{\mathrm{MN}}(x; a) : a \in M\}$$

and p is realized by $2 \in M$. So every complete φ^{MN}-type over M is realized in M, but A is not recognizable. □

References

[1] Véronique Bruyère et al. "Logic and p-recognizable sets of integers". In: *Bull. Belg. Math. Soc* 1 (1994), pp. 191–238 (cit. on p. 1).

[2] Volker Diekert and Paul Gastin. "First-order definable languages". In: *Logic and Automata: History and Perspectives*. Ed. by Jörg Flum, Erich Grädel, and Thomas Wilke. Vol. 2. Texts in Logic and Games. Amsterdam University Press, 2008, pp. 261–306. URL: http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/DG-WT08.pdf (cit. on p. 14).

[3] Anand Pillay. *Geometric Stability Theory*. Clarendon Press, 1996 (cit. on p. 15).

[4] Jean-Éric Pin. "Mathematical Foundations of Automata Theory". Feb. 1, 2016. URL: https://www.irif.univ-paris-diderot.fr/~jep/PDF/MPRI/MPRI.pdf (visited on 08/24/2016) (cit. on pp. 1, 13, 22).

[5] M.O. Rabin and D. Scott. "Finite Automata and Their Decision Problems". In: *IBM Journal of Research and Development* (Apr. 1959) (cit. on p. 20).

[6] Z. Sela. "Diophantine geometry over groups VI: the elementary theory of a free group". In: *Geometric & Functional Analysis GAFA* 16.3 (2006), pp. 707–730. ISSN: 1420-8970. DOI: 10.1007/s00039-006-0565-8. URL: http://dx.doi.org/10.1007/s00039-006-0565-8 (cit. on p. 6).

[7] Jeffrey Shallit. *A Second Course in Formal Languages and Automata Theory*. Cambridge University Press, 2008 (cit. on pp. 3–5, 22).

[8] Pierre Simon. *A Guide to NIP Theories*. Cambridge University Press, July 2015. URL: http://www.normalesup.org/~simon/NIP_guide.pdf (visited on 08/29/2016) (cit. on p. 8).

[9] Michael Sipser. *Introduction to the Theory of Computation*. Nelson Education, 2012 (cit. on p. 9).