

<https://helda.helsinki.fi>

Forgotten Islands of Regularity in Phonology

Yli-Jyrä, Anssi Mikael

MTA Research Institute for Linguistics
2019

Yli-Jyrä , A M 2019 , Forgotten Islands of Regularity in Phonology . in B Gyuris , K Mády & G Recski (eds) , K + K = 120 Papers dedicated to László Kálmán and András Kornai on the occasion of their 60th birthdays . MTA Research Institute for Linguistics , Budapest , pp. 541 - 559 .

<http://hdl.handle.net/10138/327407>

unspecified
publishedVersion

Downloaded from Helda, University of Helsinki institutional repository.

This is an electronic reprint of the original article.

This reprint may differ from the original in pagination and typographic detail.

Please cite the original version.

■ **Forgotten islands of regularity in phonology**

ANSSI YLI-JYRÄ

anssi.yli-jyra@helsinki.fi
University of Helsinki

KEYWORDS

Generative Phonology
finite-state phonology
regularity
linear-time Turing machines

ABSTRACT

Hennie presented a very general sufficient condition for regularity of Turing machines. This happened chronologically before Generative Phonology (Chomsky & Halle 1968) and the related finite-state research (Johnson 1972; Kaplan & Kay 1994). Hennie's condition lets us (1) construct a finite-state transducer from any grammar implemented by a linear-time Turing machine, and (2) to model the regularity in context-sensitive derivations. For example, the suffixation in hunspell dictionaries (Németh et al. 2004) corresponds to time-bounded two-way computations performed by a Hennie machine. Furthermore, it challenges us to look for new forgotten islands of regularity where Hennie's condition does not necessarily hold.

1. Introduction

Generative Phonology (Chomsky & Halle 1968) is a rule-based string rewriting system that has been scrutinized carefully over the years of its existence. One of the major weaknesses of the system is that it has been proven to be equivalent to Turing machines (TMs) (Chomsky 1963; Johnson 1972; Ristad 1990). As the derivations of such a machine do not necessarily terminate, the system is seriously defective and impossible to falsify. Thus, an unrestricted rewriting allowed by Generative Phonology does not make a very good scientific theory in the light of Popper (1959), see also Johnson (1972, 32).

In spite of the original shortcomings and the increased depth in the current phonological theory, the original “SPE” formalism is interesting for its own sake. First, the formalism has been employed extensively in natural language processing and descriptive linguistics. There, it has been used to express phonological generalizations based on empirical data. Second, revisiting the original formalism and its decidable subsets can produce

valuable ideas that are applicable to more ambitious theories, such as Optimality Theory (Prince & Smolensky 2004) and Harmonic Serialism (McCarthy 2000).

1.1. The well-known islands of regularity

When realistic grammar instances in Generative Phonology have been studied closely, a striking contrast between the original undecidable theory and the actual grammars has been discovered. In pioneering studies (Johnson 1972; Kaplan & Kay 1994), most practical grammars in Generative Phonology have been shown to *satisfy* a two-part condition under which they correspond to finite-state transducers:

1. *Non-self-embedding*. Directional or simultaneous context-sensitive rules whose non-contextual parts do not apply to their own output are finite-state (Kaplan & Kay 1994, 363, 365).
2. *Finite composition*. If a grammar is defined as a finite sequence of rewriting rules, each of which is a regular relation, then the grammar as a whole represents the regular relation given by their composition (*ibid.*, 364).

These observations have led to the development of algorithms for transforming restricted fragments, or *islands*, of Generative Grammar into finite-state transducers. For example, the algorithm of Mohri & Sproat (1996) constructs transducers from rules that are applied in a directed fashion. Karttunen (1995) treats various application modes and different types of context conditions. These finite-state islands in Generative Phonology have become standard textbook material (Jurafsky & Martin 2000; Beesley & Karttunen 2003), and many redesigned compilation algorithms have been proposed to pursue efficiency, flexibility and the generally correct semantics.

The literature of methods that compile individual rules into finite-state transducers suggests that regularity of phonological grammars is to be proven inductively, by using operations that preserve regularity of regular relations. But we should not overlook a more extensive picture of regularity as a property of the relation rather than as a property of the construction. Therefore, we should now start to pursue for a wider understanding of the *archipelago* of finite-state islands in Generative theories as well as in all computational models of language.

1.2. The search for further islands

Proving that the input-output relation defined by a grammar is regular is a complicated task. The known finite-state islands and the closure properties of finite-state transducers solve only the easy cases where the application order of rules is fixed and the rules can be combined under a finite composition. But if the grammar contains iterative rules, we do not have a general method that would return a non-iterative grammar. The regularity of the string relation defined by iterative rules is computationally undecidable already for context-free grammars (Stearns 1967; Greibach 1968), not to talk about Turing machines and equivalent grammar systems.

In light of this, we see that the fundamental results in finite-state Phonology (Johnson 1972; Kaplan & Kay 1994) have given us only islands, sufficient conditions where the grammars or parts of grammars are finite-state and generate regular relations. They do not exclude new conditions that can also be valuable. New conditions are, ideally, constructive and turn a formerly nonconstructive property into a method that gives a finite-state transducer.

For example, it has been obvious since Chomsky & Halle (1968) that a phonological grammar is regular when it contains only right-linear (or left-linear) rules (Chomsky 1963). The left-linear rules have the general shape $\alpha \rightarrow \beta\gamma$ where α, β, γ are symbols and β does not match any left-hand side in the grammar rules. The achievement of Johnson (1972) was to expand the default regular subset of Generative Phonology by showing that the linear and the simultaneous application of phonological rules with context conditions can also generate a regular relation.

Kaplan and Kay (1994) also discuss general situations that are usually known as cyclic derivations (Mohan 1986). For example, the word *unenforceable* has the recursive structure $[un[[en[force]]able]]$. Here the phonological rules are applied first to the innermost part, *force*. Then the innermost brackets are removed and the application is repeated until no brackets are left. Kaplan and Kay point out that “there may be restrictions on the mode of reapplication that limit the formal power of the [cyclic] grammar...”. However, Kaplan and Kay (1994, 365) seem to think that these restrictions are analogous to context-free grammars with only right- or left-linear rules.

Besides context-free grammars with only right- or left-linear rules, there are also self-embedding grammars that generate regular languages. For example, the context-free grammar $S \rightarrow aS \mid Tb; \quad T \rightarrow Tb \mid c$ generates the regular language a^*cb^* and the context-sensitive grammar

$S \rightarrow aS$; $aS \rightarrow abT$; $bT \rightarrow cbT \mid c$ generates the regular language a^*c . In these examples, the grammars look simple but are not immediately regular on the basis of the shape of their rules.

In the sequel, section 2 presents a concrete example of a very simple context-sensitive formalism whose conversion to a finite-state equivalent grammar is tricky. In section 3, the reader is familiarized with a one-tape Turing machine and Hennie's sufficient condition for regularity. The paper closes with remarks in section 4.

2. Safe unbounded composition

We will now give an example of a grammar whose regularity is not obvious on the basis of the standard conditions.

2.1. hunspell

Our example of a non-classical finite-state grammar is the **hunspell** formalism (Németh et al. 2004) that represents a stage in the development of spell checking algorithms. We only discuss its suffix rules and ignore many details of the formalism.

The **hunspell** formalism is used to inflect and derive word forms by a combination of continuation classes, truncation and appending. The formalism resembles the Item and Process morphology (Hockett 1954) and Lexical Phonology (Mohanalan 1986). The formalism involves `.dic` and `.aff` files that specify the initial word forms and the steps to produce other word forms:

- (1) `.dic: glossy/T`
`.aff: SFX T y iest Cy`

The word form *glossiest* is the combination of an input word *glossy*, having the continuation class T, and a suffix rule (marked with SFX). The suffix rule, for the continuation class T (T in the 2nd column), states that the last vowel *-y* (y in the 3rd column) is replaced with *-iest* (*iest* in the 4th column) if preceded by a consonant and the vowel *-y* (condition Cy in the 5th column).

When the **hunspell** dictionary formalism is interpreted as a rewriting system, we see that the derivation `glossyT` \Rightarrow `glossiest#` is described with a Generative Phonological rule (2):

(2) Superlative Formation:

$$[y]T \rightarrow [i][e][s][t][\#] / C _$$

While the shape of such a rule is context-sensitive, it is not difficult to see that this rule can be implemented with a non-deterministic finite-state transducer. Furthermore, the suffix rules seem to be applied out from the stem, at the right boundary of the string. However, such similarity with right-linear grammars is only partial and does not imply that it would be easy to compile the whole dictionary into a finite-state transducer. There are two reasons:

- The rules do not only rewrite the continuation classes but they may also back up and rewrite the phonological content produced earlier, requiring, thus, *two-way* movements.
- The rules are *non-monotonic*: they can expand and shorten the input.


In order to analyse what actually happens, we need to construct a model that shows how the dictionary form is processed by the affix rules.

2.2. Automaton models

Now we analyse *hunspell* by viewing the derivation steps of its word formation as a process that corresponds to a computation by a particular TM.

The general definition of a Turing machine is assumed to be familiar to the reader. In short, it is a combination of a finite-state automaton and a rewritable two-way working tape that is initialized with the input string of length *n*. The machine is allowed to append new letters arbitrarily to the input string; thus the working tape is infinite. A TM can also have auxiliary tapes, but we will restrict ourselves to one-tape TMs.

If we implement the derivation by the moves of a *non-deterministic* one-tape TM, we obtain a machine that sweeps the working tape three times in a row. During the first pass (3.1), the machine recognizes the stem (*glossy*) and its continuation class (T), then rewinds the tape (3.2) to the beginning of the string and non-deterministically replaces the substring *syT\$\$\$* with the substring *siest#* during the final pass (3.3):




Since regular relations are closed under composition, a finite number of similar suffix rules could be applied in a row and the lexicon would still be regular. In this way, k suffix positions of morphology could be treated. The corresponding non-deterministic TM would rewind the tape k times to the beginning. Thus, the total time complexity is in $O(nk)$ when the string on the tape occupies at most n tape squares. Thus, the non-deterministic TM implementation of a finite composition has linear time complexity.


In a more general situation, one does not want to specify the maximum number of suffixes explicitly. One reason can be that, in some languages, the suffixes can be added recursively after one other. For example, a Turkish word can, in principle, have an arbitrary number of suffixes although only some of the combinations are interpretable. Another example involves Old Georgian where the nouns can theoretically have unlimited number of case-number markers (Michaelis & Kracht 1997). Finally, orthographic compounding of many languages can involve several stems and alternating bound morphemes. For example, the Swedish word (with our morpheme boundaries) *Spår-vagn-s-aktie-bolag-s-sken-smut-s-skjut-are-fack-förening-s-personal-beklädnad-s-magasin-s-förråd-s-förvaltar-en-s* contains 14 stems and, in addition, several bound morphemes. Similar and much longer examples can be found in other languages. E.g., a 431-letter word appears in the Sanskrit literature. Thus, it is hard to argue that the number of rule applications has a finite upper bound in general.

The unbounded number of applications of suffix rules breaks the two principles: *non-self-embedding rules* and *finite composition*. Moreover, the implementation based on a rewinding TM would spend $O(n^2)$ time to produce the derived string through the back and forth sweeps that simulate the composition steps.

We can, however, improve the TM implementation by optimizing its moves. Instead of rewinding the tape completely, the improved computation strategy (4) just backs up until it has tested the precondition. In our example, the precondition is just the suffix $[C][y][T]$:



With this change, long words are produced in a zigzag style (5) where every rule application may back up some letters.



Since the union of the affix-rules is applied repeatedly to its own output, the standard two-part regularity condition of phonological grammars does not apply. However, as long as the derivation deletes and appends new material only at the right end of the string, the resulting process is linear and, intuitively, a regular grammar. In addition, the moves taken by the TM can now be deterministic because the machine does not completely rewind the tape at any point but always makes relative moves that allow it to remember its previous position.

2.3. Linear encoding

Although the grammar represented by a `hunspell` lexicon does not satisfy the classical two-part condition of finite-state phonology, it is equivalent to a finite-state transducer when restricted to the suffix rules.

There are now some methods to compile `hunspell` lexicons to finite-state transducers. Early experiments on compilation are due to György Gyepesi (p.c., 2007) and others in Budapest. The author developed his solution (Yli-Jyrä 2009) using a variant of Two-Level Morphology (Koskeniemi 1983). This method viewed the lexicon as a collection of constraints that described linearly encoded backing up and suffixation in derivations. The method included an efficient one-shot compilation algorithm to compile and intersect several hundreds of thousands of lexical context restriction rules in parallel as if the lexical continuations (morphotaxis) were phonological constraints. A similar method, finally implemented by his colleagues, Pirinen and Linden (2010), separated the lexical continuations from the phonological changes at morpheme boundaries and used a three-step approach where the final step composed the lexicon with the phonology. A separate compiler for lexical continuations was used and a two-level grammar described the phonological realization of morphemes in different contexts.

The key to understanding the method of Yli-Jyrä (2009) is that the computations of the TM are encoded as a linear string (6.2). This string is produced if the derivation actually reaches the final continuation class # whereas infinite loops do not correspond to an output string.

(6) gloss y (6.1)

<#><Root>glossδy<T>δ⁻¹iest<#> (6.2)

gloss iest. (6.3)

Whenever the machine overwrites its own output, the overwritten part (already on the left from the current position) is marked as deleted. The deleted material is put between <D> and <-D>, two symbols abbreviated now as δ and δ⁻¹, respectively. The occurrence of these optional lexical symbols is enforced in deleting contexts but banned otherwise, corresponding to the surface realizations with and without contractions. For example, the segment y in (6.2) is surrounded by a pair of δ and δ⁻¹ because it is cancelled by -iest, the next *hunspell* affix.

The derivation of the combination of *glossy* and -iest is encoded as string (6.2). This internal string is then mapped to the output string (6.3) by a transducer that deletes the markers and the material enclosed between each pair of δ and δ⁻¹. The computation can also be mapped to the dictionary form (6.1) by removing the markers and the material that belongs to the affixes.

An interesting part in the method is that the underlying derivation encodes a computation of a bounded Turing machine (6.2). This string is produced with context-restriction rules introduced in two-level phonology (Koskeniemi 1983). Since the output deletions are taken care of by the simple transducer between (6.2) and (6.3), it is sufficient to describe only one representation level. The three rules in (7) describe where the root (Root) of the lexicon is visited, where a continuation class T is reached, and how the next (in fact final) continuation class is reached after a consonant, a cancelled y and new material corresponding to lexical -iest.

(7) <Root> => <#> _ ; # the root symbol
 <T> => <#>δ*<Root>δ*gδ*lδ*oδ*sδ*sδ*yδ* _ ; # glossy/T
 <#> => Cδy<T>δ⁻¹δ*iδ*eδ*sδ*tδ*<#>δ* _ ; # SFX T y iest Cy

The one-level representation of the underlying derivation works immediately in cases where the successive *deletions* are disjoint from each other. It can also be extended to cases where the deleted parts are nested:

$$\begin{array}{c}
 \text{abcef}<\text{A}> \quad \text{gh}<\text{B}> \quad \text{hij}<\text{C}> \\
 \underbrace{\text{abc}\delta\text{e}\delta\text{f}<\text{A}> \delta^{-1}\text{gh}<\text{B}> \delta^{-1}\text{hij}<\text{C}>} \\
 \underbrace{\hspace{1.5cm}} \text{truncate f before gh}<\text{B}> \\
 \underbrace{\hspace{2.5cm}} \text{truncate egh before hij}<\text{C}>
 \end{array}
 \tag{8}$$

The main functional difference between the methods described by Yli-Jyrä (2009) and Pirinen & Lindén (2010) is in the way they treat non-disjoint deletions. While the former method encodes the sequence of derivation steps as one string, the latter encodes the lexical morpheme sequences on one string and then the contracted sequences on the other level. The latter method describes the contractions at morpheme boundaries via two-level rules that constrain the way in which the underlying phonemes of a morpheme are realized in the adjacency of various affixes. In this approach, a contraction corresponds to zero realization.

$$\begin{array}{c}
 \text{abcef}<\text{A}> \quad \text{gh}<\text{B}> \quad \text{hij}<\text{C}> \\
 \underbrace{\text{abc e f}<\text{A}> \text{gh}<\text{B}> \text{hij}<\text{C}>} \\
 \text{abc 0 0 0 0 0 hij 0} \\
 \underbrace{\hspace{1.5cm}} \text{truncate f before gh}<\text{B}> \\
 \underbrace{\hspace{2.5cm}} \text{truncate e f g h before hij}<\text{C}>
 \end{array}
 \tag{9}$$

Since the truncations in this representation (9) are specified in parallel rather than one after another, the semantics of the variant (Pirinen & Lindén 2010) deviates slightly from the original method (Yli-Jyrä 2009).

In particular, note that the addition of the suffix *hij*<C> in (8) and (9) requires different suffixation rules as the truncations behave differently. The rule applied in (9) must truncate more symbols. This semantic difference between the two methods can be compensated with an additional pre-processing step that expands the set of suffixation rules. During this step, a suffix rule that completely cancels the previous affix is replaced with a suffix rule that is applied before the completely cancelled affix. However, for most *hunspell* lexicons, the cancellation is restricted to the most recent suffix, which means that the preprocessing step can be heuristically ignored.

3. The loosest sufficient condition

In the previous section, we related the `hunspell` derivations to one-tape TMs. One reason to do so was that the regularity of one-tape TMs is an old and carefully studied, well-understood problem.

In this section, we first relate one-tape Turing machines with transducers (§3.1–3.2). Then we study bounded one-tape TMs that implement regular relations (§3.3). We will use the bounded one-tape TMs to give a new proof for the two-part regularity condition (§3.4) and to find a more general condition for a finite-state subset of Generative Phonology (§3.5). Finally, we observe (§3.6) that even this condition does not cover all natural finite-state grammars.

3.1. One-tape TMs as transducers

Usually one-tape TMs and Hennie machines are viewed as language recognizers. Since it is not possible to construct a Hennie machine with two readable tapes (Hennie 1965), the connection between Hennie machines and one-way two-tape finite-state transducers is not obvious from the beginning. In fact, most of the relevant literature discusses Hennie machines as if they were equivalent to one-tape finite-state automata only.

As a notable exception, Engelfriet and Hoogeboom (2001) connect Hennie machines to two-way two-tape finite-state machines. These machines are not allowed to read their output tape, but they are more powerful than ordinary finite-state transducers.

Our way to view one-tape Turing machines as transducers requires only the input tape with both reading and writing. As the machine modifies the contents of the input tape during its computations, the input tape will be occupied with an output string when the machine halts. Thus, every one-tape TM recognizes three sets:

- *the set of input strings* that occur as the initial content of the working tape in an accepting computation,
- *the set of output strings* that occur as the final content of the working tape in an accepting computation,
- *the relation consisting of the input-output string pairs* where the first string is the initial content and the second string is the final content of the working tape in an accepting computation.

Given the last definition, every one-tape TM can be viewed as a recognizer of a binary relation.

3.2. Finite-state transducers as one-tape TMs

It is immediate that one-way finite-state transducers are equivalent to one-tape Turing machines: one-way (non)deterministic finite-state transducers are a special case of two-way (non)deterministic finite-state transducers, and these are a special case of (non)deterministic two-tape Turing machines that are equivalent to deterministic one-tape Turing machines.

If we restrict ourselves to finite-state transducers whose output preserves the length of their input, we can view these transducers as one-tape finite automata with a letter-pair alphabet (Kaplan & Kay 1994). This gives an even more direct link from finite-state transducers to one-tape Turing machines.

The restriction to these *letter transducers* is not a serious restriction if we assume that the necessary 0's are introduced non-deterministically to the input string by an inverse homomorphic mapping $h^{-1} : \Sigma^* \rightarrow (\Sigma \cup \{0\})^*$ that preserves the alphabet Σ . The 0's are then removed from the output string by the homomorphism $h : (\Sigma \cup \{0\})^* \rightarrow \Sigma^*$. In addition, we assume that all states of the transducer have a self-loop on the letter pair (0, 0).

Let R_1 and R_2 be regular relations recognized by unrestricted finite-state transducers, and let R'_1 and R'_2 be the corresponding same-length relations recognized by the letter transducers with the self-loops. Now we have the equation:

$$R_1 \circ R_2 = h^{-1} \circ R'_1 \circ R'_2 \circ h.$$

It is now obvious that the same length relations R'_1 , R'_2 and even $R'_1 \circ R'_2$ can be implemented as non-deterministic one-tape TMs that recognize the relations in $O(n)$ time by transforming the initial content of the tape to the final content of the same tape.

3.3. TMs running in $O(n)$ time

The most ground-breaking regularity condition for one-tape TMs is due to Hennie. Hennie's result is the converse to the fact that every one-way deterministic finite automaton is a deterministic TM. The machines considered by Hennie do not only include all one-way deterministic finite automata and letter transducers but they also extend them in two particular ways:

(1) the one-tape TMs can move back and forth on the tape, (2) they can overwrite the contents of the squares of the tape several times.

Hennie (1965) showed that a deterministic one-tape TM is equivalent to a finite automaton if it runs in $O(n)$. The results of Hennie have been extended, by Tadaki et al. (2010), to linear-time non-deterministic one-tape TMs whose $O(n)$ time bound holds for all accepting computations. A deterministic and non-deterministic linear-time one-tape TM are called a *Hennie machine* and a *non-deterministic Hennie machine*, respectively. Both of these one-tape machines recognize a regular relation on the basis of section 3.1.

Hennie analysed the expressive power of one-tape machines using the concept of *crossing sequence* (aka *schema*) (Rabin 1963; Trakhtenbrot 1964; Hopcroft & Ullman 1979; Birget 1996) that is strongly related to *visiting sequences* (Fischer 1969). This concept is a powerful tool in the analysis of the behaviour of two-way automata and one-tape TMs. It refers to the sequence of target states s_1, s_2, \dots visited by a TM when its pointer crosses the boundary between a pair of adjacent tape squares. States s_1, s_3, \dots are reached when the pointer moves forward and states s_2, s_4, \dots are reached when the pointer moves backwards. Figure 1 shows how states are visited during a computation and how a crossing sequence is defined.

#	W	O	R	D	I	N	G	#	#	...
q_0	q_1	q_2	q_3	q_4						
	q_7	q_6		q_5	\leftarrow					
	\hookrightarrow	q_8	q_9	q_{10}	q_{11}	\dots				

Figure 1: The crossing sequence between the 3rd and the 4th squares is $(s_1, s_2, s_3) = (q_3, q_6, q_9)$

Every Hennie machine satisfies, by definition, the property that the length of its crossing sequences is bounded by an integer (Hennie 1965). Since the finiteness of crossing sequences implies that the TM is equivalent to a finite-state automaton, this bound lets us construct an equivalent finite-state device. The good news is that for all Hennie machines, a bounding constant k is computable (Kobayashi 1985; Tadaki et al. 2010).

Průša (2014) showed that, if a deterministic Hennie machine recognizing the input language has m states and n working symbols, we can construct a minimal deterministic finite automaton that has $2^{2^{O(n \log m)}}$ states. Thus, Hennie machines recognizing the input are much more succinct than

the equivalent minimal deterministic automaton. Obviously, a deterministic letter transducer constructed from a Hennie machine is not smaller than the minimal automaton recognizing only the input language aka the domain of the transducer.

3.4. Completeness with respect to prior art

Now we can prove that Hennie machines can be used, on the one hand, to build a finite-state subset of Generative Phonology using the two-part condition of non-self-application and finite composition, and, on the other, to obtain compilation methods for such grammars that previously required specialized encoding and a compilation algorithm.

Theorem 1. *Composition of regular relations is a regular relation.*

Proof. Let R_1 and R_2 be regular relations and R'_1 and R'_2 the respective same-length regular relations. Then there are, respectively, two non-deterministic Hennie machines T_1 and T_2 that recognize R'_1 and R'_2 . The composition $R'_1 \circ R'_2$ is then computed by a Hennie machine that first runs like T'_1 , then rewinds the tape and runs like T'_2 . Since the combined machine preserves the string length, it is equivalent to an epsilon-free finite-state transducer that recognizes the relation $R'_1 \circ R'_2$. The composition $h^{-1} \circ R'_1 \circ R'_2 \circ h$ is then equivalent to the regular relation $R_1 \circ R_2$. \square

Theorem 2. *The non-self-embedding application of rule of the form $\alpha \rightarrow \beta/\lambda_ \rho$ corresponds to a regular relation.*

Proof. Extend the original tape alphabet so that each square contains the input letter and a Boolean vector indicating the validity of left and right context conditions of the simultaneous rules. Let M_L (M_R) be a deterministic (co-deterministic) pattern matching automaton. The state computed by this automaton indicates, for each string position, the type of the prefix (suffix) of the position. Modify this pattern matching automaton by adding self-loops on 0's. Then transform the automaton into a finite-state transducer M'_L (M'_R) in such a way that each transition adds the information on the occurring left (right) contexts to the Boolean vectors of each square. This epsilon-free transducer is a Hennie machine. The composition $M'_L \circ M'_R$ is then a Hennie machine that marks the occurring contexts at all squares.

As a pre-processing step, make the length of each left-hand side α and the respective right-hand side β identical by padding the shorter with 0's.

In addition, add synchronous 0's freely to both. In this way we obtain a letter transducer that recognizes a 0-padded representation of the regular relation $\alpha \times \beta$.

Define a Hennie machine M_1 that sweeps the string (containing 0's) from left to right and non-deterministically overwrites ranges of squares that contain some left-hand-side string α with a corresponding right-hand-side string β when the first and the last square in the input range indicate the presence of the required left and right context, respectively. Define also a Hennie machine M_2 that removes the Boolean context vectors from the tape squares. Now the composition $M'_L \circ M'_R \circ M_1 \circ M_2$ is recognized by a Hennie machine M . Then $h^{-1} \circ M \circ h$ is equivalent to a non-deterministic finite-state transducer that captures the semantics of the rule. \square

We have now used Hennie machines to show that simultaneous *non-overlapping* rules are regular and that a *finite composition* of regular rules preserves regularity. Other application modes of regular grammars are discussed in Johnson (1972); Kaplan & Kay (1994). The regularity of these application modes can be proven similarly.

To conclude our argument, we show that Hennie machines actually help us to compile `hunspell` dictionaries without special encodings.

Theorem 3. *The iterated application of monotonic suffix rules of a `hunspell` grammar makes a regular relation.*

Proof. Every suffix rule corresponds to a Hennie machine that backs up checking its context condition and then writes the non-truncated context and the new suffix (4.2–4.3). The union of such Hennie machines is a non-deterministic Hennie machine M . The closure M^* is a TM that applies suffix rules iteratively. As the suffix rules increase the length of the string monotonically, the closure M^* has a finite bound for the crossing sequences and recognizes a regular relation. \square

3.5. The bound that cannot be improved

The Borodin-Trakhtenbrot Gap Theorem (Trakhtenbrot 1964) states that expanded resources do not always expand the set of computable functions. In other words, it is possible that the regularity of a TM holds even if the $O(n)$ is made slightly looser. A less tight time bound is now expressed with the small-o notation: $t(n) \in o(f(n))$ means that the upper bound $f(n)$ grows much faster than the running time $t(n)$ when n tends to infinity: $\lim_{n \rightarrow \infty} t(n)/f(n) = 0$.

As an application of the Gap Theorem, Hartmanis (1968) and Trakhtenbrot (1964) showed independently that the time resource of finite-state equivalent deterministic one-tape TMs can be extended from $O(n)$ to $o(n \log n)$. This bound is tight: regularity is algorithmically unsolvable for any bound that exceeds $n + 1$ in $\Omega(n \log n)$ (Gajser 2015). The extended time bound has been generalized to non-deterministic one-tape TMs by Tadaki et al. (2010). These extensions of Hennie's core result give us a new sufficient condition for the regularity of Generative Phonology.

Theorem 4. *A generative phonological grammar is regular if its one-tape TM implementation runs in $o(n \log n)$ time.*

Let M be a one-tape TM implementation of a Generative phonological grammar. The finiteness of the crossing sequences of a given TM is, in general, undecidable (Průša 2014), but there is a reasonably good decision procedure: to test if M is equivalent to a finite-state transducer, we can pick a function $t(n)$ that is in $o(n \log n)$ and test if M actually runs in $t(n)$. Interestingly, Gajser (2015) showed that for any reasonable function $t(n)$, we can decide whether a TM M runs in $t(n)$. If a TM then runs in $t(n)$, it actually runs in $O(n)$ (Pighizzini 2009). Thus, the new one-sided condition for regularity of the phonological grammar has a sound approximate solution.

3.6. The existence of non-Hennie finite-state grammars

If the suffix rules are non-monotonic and can shorten their inputs, the TM can produce the same configuration again and again and produce arbitrarily long crossing sequences. The repetition may happen either a finite or an infinite number of times. Interestingly, the specialized compilation method (Yli-Jyrä 2009) handles both cases correctly, whereas we fail to get a Hennie machine if the suffix rules are non-monotonic.

Non-monotonic suffix rules are an example of a situation where the TM is equivalent to a Hennie machine that restricts the length of the crossing sequences. The bad news is that we do not know when we have a correct Hennie machine: it is difficult to find such bound k for the length of crossing sequences that a given TM preserves its semantics when longer crossing sequences are abandoned. Since a sufficient bound k is such that the semantics of the restricted TM does not change although we allow longer crossing-sequences, there are reasonable ways to probe possible values of k , but such probing is still heuristic.

The difficulty of non-monotonic grammars indicates that although we now have a more general condition for those Generative phonological grammars that are equivalent to a finite-state transducer, a specialized compilation algorithm may still encode infinite loops in a way that seems to be beyond the Hennie condition.

4. Conclusions

It is historically interesting that Hennie's regularity condition dates back to the year 1965, that is, even before Chomsky & Halle (1968).

No decision procedure for the classical two-part condition (Johnson 1972; Kaplan & Kay 1994), is known. Compared to this situation, it is remarkable that the new sufficient condition has several advantages:

- The new regularity condition has approximations that are decidable (Gajser 2015).
- The equivalent finite-state transducer can be constructed from a Hennie machine (Hennie 1965).
- The Hennie machines are extremely succinct compared to finite-state machines (Průša 2014).
- Hennie machines seem to provide a more general framework for proving regularity of phonological grammars than the arguments based on bimachine construction (Johnson 1972) or non-self-embedding grammars (Kaplan & Kay 1994).

There are many interesting questions that could be studied in the future. Here are some:

1. Despite the advantages of Hennie machines, the author is not aware of any finite-state library that would be based on Hennie machines. *Would it be possible to develop a finite-state library that would use Hennie machines to represent regular relations more compactly?*
2. There does not seem to be much work that would link Hennie machines and two-way finite-state transducers to *minicomplexity*, the computational complexity of two-way finite automata, that has recently obtained attention in automata theory (Kapoutsis 2012). *Could some of the related results be extended to Hennie machines?*
3. If a Hennie machine is used to implement a weighted rule system, the machine must be constructed more carefully than what we have done now: the 0-loops create new paths that make the computation

of string weights tricky. *Can we introduce weighted Hennie machines and relate them to weighted automata?*

4. We would like to understand why some natural finite-state grammars, like nonmonotonic `hunspell` grammars, are finite-state, although their crossing sequences seem to have no finite bound. *Are there thus other natural islands of regularity we should know about?*

There are several potential applications for Hennie machines in Natural Language Processing. We have already demonstrated that Hennie machines have applications in phonology and morphology (Yli-Jyrä 2009). Weighted Hennie machines may be applied to OCR that is based on weighted context-dependent correction rules (Drobac et al. 2017). Furthermore, non-monotonic Sequential Constraint Grammar is computationally undecidable but has restrictions that have Hennie machine characterizations (Yli-Jyrä 2017). The search for Generative dependency grammars that produce non-projective trees is an area that may also benefit from the concepts of crossing sequences and Hennie machines (Nederhof & Yli-Jyrä 2017).

Acknowledgements

Writing this article was supported by the author's fellowship #270354 from the Academy of Finland.

References

- Beesley, K. R. and L. Karttunen. 2003. Finite state morphology. Stanford, CA: CSLI Publications.
- Birget, J.-C. 1996. Two-way automata and length-preserving homomorphisms. *Mathematical Systems Theory* 29. 191–226.
- Chomsky, N. 1963. Formal properties of grammars. In R. Luce, R. Bush and E. Galanter (eds.) *Handbook of mathematical psychology*. New York: John Wiley and Sons. 323–418.
- Chomsky, N. and M. Halle. 1968. *The sound pattern of English*. New York: Harper & Row.
- Drobac, S., P. Kauppinen and K. Lindén. 2017. OCR and post-correction of historical Finnish texts. In *Proceedings of the 21st Nordic Conference of Computational Linguistics*. Gothenburg: Linköping University Electronic Press, 70–76.
<http://www.ep.liu.se/ecp/131/009/ecp17131009.pdf>
- Engelfriet, J. and H. J. Hoogeboom. 2001. MSO definable string transductions and two-way finite-state transducers. *ACM Transactions on Computational Logic* 2. 216–254.

- Fischer, M. J. 1969. Two characterizations of the context-sensitive languages. In *IEEE Conference Record of 10th Annual Symposium on Switching and Automata Theory*. 149–156. <http://ieeexplore.ieee.org/document/4569611/>
- Gajser, D. 2015. Verifying time complexity of Turing machines. *Theoretical Computer Science* 600. 86–97.
- Greibach, S. 1968. A note on undecidable properties of formal languages. *Mathematical Systems Theory* 2. 1–6.
- Hartmanis, J. 1968. Computational complexity of one-tape Turing machine computations. *Journal of the ACM* 15. 325–339.
- Hennie, F. C. 1965. One-tape, off-line Turing machine computations. *Information and Control* 8. 553–578.
- Hockett, C. F. 1954. Two models of grammatical description. *Word* 10. 210–231.
- Hopcroft, J. E. and J. D. Ullman. 1979. *Introduction to automata theory, languages and computation*. Reading, MA: Addison-Wesley.
- Johnson, C. D. 1972. *Formal aspects of phonological description*. The Hague: Mouton.
- Jurafsky, D. and J. H. Martin (eds.). 2000. *Speech and language processing: An introduction to natural language processing, computational linguistics, and speech processing*. Englewood Cliffs, NJ: Prentice-Hall.
- Kaplan, R. M. and M. Kay. 1994. Regular models of phonological rule systems. *Computational Linguistics* 20. 331–378.
- Kapoutsis, C. A. 2012. Minicomplexity. In M. Kutrib, N. Moreira and R. Reis (eds.) *Descriptive Complexity of Formal Systems: 14th International Workshop, DCFS 2012, Braga, Portugal, July 23–25, 2012. Proceedings*. Berlin & New York: Springer. 20–42.
- Karttunen, L. 1995. The replace operator. In *33rd Annual Meeting of the Association for Computational Linguistics*. 16–23. <https://dl.acm.org/citation.cfm?id=981658.981661>
- Kobayashi, K. 1985. On the structure of one-tape nondeterministic Turing machine time hierarchy. *Theoretical Computer Science* 40. 175–193.
- Koskenniemi, K. 1983. *Two-level morphology: A general computational model for word-form recognition and production*. Helsinki: University of Helsinki.
- McCarthy, J. J. 2000. Harmonic serialism and parallelism. In M. Hirotani, A. Coetzee, N. Hall and J.-Y. Kim (eds.) *Proceedings of NELS 30*. Amherst: University of Massachusetts. 501–524.
- Michaelis, J. and M. Kracht. 1997. Semilinearity as a syntactic invariant? *Logical Aspects of Computational Linguistics* 1328. 329–345. <https://doi.org/10.1007/BFb0052165>
- Mohanan, K. P. 1986. *The theory of lexical phonology*. Dordrecht: D. Reidel.
- Mohri, M. and R. Sproat. 1996. An efficient compiler for weighted rewrite rules. In *34th Annual Meeting of the Association for Computational Linguistics*. 231–238. <http://aclweb.org/anthology/P96-1031>
- Nederhof, M.-J. and A. Yli-Jyrä. 2017. A derivational model of discontinuous parsing. In *Language and Automata Theory and Applications – 11th International Conference, LATA 2017, Umeå, Sweden, March 6–9, 2017. Proceedings*. 299–310. https://doi.org/10.1007/978-3-319-53733-7_22
- Németh, L., V. Trón, P. Halácsy, A. Kornai, A. Rung and I. Szakadát. 2004. Leveraging the open source ispell codebase for minority language analysis. In J. Carson-Berndsen (ed.) *Proceedings of the SALT MIL Workshop at LREC 2004*. 56–59.

- Pighizzini, G. 2009. Nondeterministic one-tape off-line Turing machines and their time complexity. *Journal of Automata, Languages and Combinatorics* 14. 107–124.
- Pirinen, T. A. and K. Lindén. 2010. Creating and weighting Hunspell dictionaries as finite-state automata. *Investigationes Linguisticae* 21. 1–16.
- Popper, K. R. 1959. *The logic of scientific discovery* (Second edition). London: Hutchinson & Co.
- Prince, A. and P. Smolensky. 2004. *Optimality Theory: Constraint interaction in Generative Grammar*. Malden, MA & Oxford: Blackwell.
- Průša, D. 2014. Weight-reducing Hennie machines and their descriptonal complexity. In A.-H. Dediu, C. Martín-Vide, J.-L. Sierra-Rodríguez and B. Truthe (eds.) *Language and Automata Theory and Applications: 8th International Conference, LATA 2014, Madrid, Spain, March 10–14, 2014. Proceedings*. Cham: Springer. 553–564.
- Rabin, M. O. 1963. Real time computation. *Israel Journal of Mathematics* 1. 203–211.
- Ristad, E. S. 1990. Computational structure of generative phonology and its relation to language comprehension. In *Proceedings of the 28th Annual Meeting on Association for Computational Linguistics*. 235–242. <http://dx.doi.org/10.3115/981823.981853>
- Stearns, R. E. 1967. A regularity test for pushdown machines. *Information and Control* 11. 323–340.
- Tadaki, K., T. Yamakami and J. C. H. Lin. 2010. Theory of one-tape linear-time Turing machines. *Theoretical Computer Science* 411. 22–43.
- Trakhtenbrot, B. A. 1964. Turing computations with logarithmic delay (in Russian). *Algebra i Logika* 3. 33–48. English translation in U. of California Computing Center, Tech. Report. No. 5, Berkeley, CA, 1966.
- Yli-Jyrä, A. 2009. An efficient double complementation algorithm for superposition-based finite-state morphology. In K. Jokinen and E. Bick (eds.) *NODALIDA 2009 Conference Proceedings*. 206–213. <http://dspace.ut.ee/handle/10062/9765>
- Yli-Jyrä, A. 2017. The power of constraint grammars revisited. In *Proceedings of the NoDaLiDa 2017 Workshop on Constraint Grammar – Methods, Tools and Applications, 22 May 2017, Gothenburg, Sweden*. 23–31. <http://www.ep.liu.se/ecp/140/005/ecp17140005.pdf>

