


Decidability of Cutpoint Isolation for Probabilistic Finite Automata on Letter-Bounded Inputs

Paul C. Bell 

Department of Computer Science, Liverpool John Moores University, UK
p.c.bell@ljmu.ac.uk

Pavel Semukhin 

Department of Computer Science, University of Oxford, UK
pavel.semukhin@cs.ox.ac.uk

Abstract

We show the surprising result that the cutpoint isolation problem is decidable for probabilistic finite automata where input words are taken from a letter-bounded context-free language. A context-free language \mathcal{L} is letter-bounded when $\mathcal{L} \subseteq a_1^* a_2^* \cdots a_\ell^*$ for some finite $\ell > 0$ where each letter is distinct. A cutpoint is isolated when it cannot be approached arbitrarily closely. The decidability of this problem is in marked contrast to the situation for the (strict) emptiness problem for PFA which is undecidable under the even more severe restrictions of PFA with polynomial ambiguity, commutative matrices and input over a letter-bounded language as well as to the injectivity problem which is undecidable for PFA over letter-bounded languages. We provide a constructive nondeterministic algorithm to solve the cutpoint isolation problem, which holds even when the PFA is exponentially ambiguous. We also show that the problem is at least NP-hard and use our decision procedure to solve several related problems.

2012 ACM Subject Classification Theory of computation \rightarrow Formal languages and automata theory; Theory of computation \rightarrow Computability; Theory of computation \rightarrow Probabilistic computation

Keywords and phrases Probabilistic finite automata, cutpoint isolation problem, letter-bounded context-free languages

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2020.22

Related Version A full version of the paper is available at <https://arxiv.org/abs/2002.07660>.

Acknowledgements We thank the referees for their careful reading of this manuscript and helpful comments.

1 Introduction

Probabilistic finite automata (PFA) are an extension of classical nondeterministic finite automata (NFA) where transitions, for each state and letter, are represented as probability distributions. The PFA model was first introduced by Rabin [23].

There are a variety of classical problems for PFA. Let \mathcal{P} denote a PFA, Σ an alphabet and $\lambda \in [0, 1]$ a probability. The acceptance probability of \mathcal{P} on a word $w \in \Sigma^*$ is denoted $f_{\mathcal{P}}(w)$. A central question is (strict) emptiness of cutpoint languages: does there exist a finite input word w for which $f_{\mathcal{P}}(w) \geq \lambda$ (or $f_{\mathcal{P}}(w) > \lambda$ for strict emptiness). Another important problem is that of *cutpoint isolation* – to determine if λ can be approached arbitrarily closely, i.e., for each $\epsilon > 0$, does there exist a word $w \in \Sigma$ such that $|f_{\mathcal{P}}(w) - \lambda| < \epsilon$ (or the converse, does there exist $\delta > 0$ such that $|f_{\mathcal{P}}(w) - \lambda| \geq \delta$ for all $w \in \Sigma^*$)? The *value-1* problem is a special case of the cutpoint isolation when $\lambda = 1$ [13]. In the *injectivity problem* we must determine if $f_{\mathcal{P}}(w)$ is injective (i.e. do there exist two distinct words with the same acceptance probability?) In the λ -*probability problem* we must determine if there exists $w \in \Sigma^*$ such that $f_{\mathcal{P}}(w) = \lambda$.



© Paul C. Bell and Pavel Semukhin;
licensed under Creative Commons License CC-BY
31st International Conference on Concurrency Theory (CONCUR 2020).

Editors: Igor Konnov and Laura Kovács; Article No. 22; pp. 22:1–22:16

Leibniz International Proceedings in Informatics



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

The emptiness problem is undecidable for rational matrices [22], even over a binary alphabet when the PFA has dimension 46 [6], later improved to dimension 25 [18]. The injectivity problem for PFA is undecidable [2], even for polynomially ambiguous PFA [4].

The main focus of this paper is the cutpoint isolation problem. The authors of [5] show that the problem of determining if a given cutpoint is isolated (resp. if a PFA has any isolated cutpoint) is undecidable and this was shown to hold even for PFA with 420 (resp. 2354) states over a binary alphabet [6]. The cutpoint isolation problem, in the special case where $\lambda = 1$ (the value-1 problem), is also known to be undecidable [13]. The problem is especially interesting given the seminal result of Rabin that if a cutpoint λ is isolated, then the cutpoint language associated with λ is necessarily regular [23].

Most problems are undecidable for PFA and there exist very few algorithmic solutions [13]. Various classes of restrictions on PFA are possible, related to the number of states, the alphabet size and whether one defines the PFA over the algebraic reals or the rationals. Recent work has studied PFA with finite, polynomial or exponential ambiguity (in terms of the underlying NFA) [10], PFA defined for restricted input words (e.g. those coming from bounded or letter-bounded languages) [2, 3], commutative PFA, where all transition matrices commute, for which cutpoint languages and non-free languages generated by such automata become commutative [4] or other structural restrictions on the PFA such as $\#$ -acyclic automata, for which some problems become decidable [13], including the value-1 problem. Such $\#$ -acyclic automata impose a restriction on the structure of the PFA (as we shall see, we only restrict the input words).

A natural restriction on PFA was studied in [3], where input words of the PFA are restricted to be from a letter-bounded language (also known as a *letter-monotonic language*) of the form $\mathcal{L} = a_1^* a_2^* \cdots a_\ell^*$ with distinct letters $a_i \in \Sigma$. This is analogous to a 1.5-way PFA, whose read head may “stay put” on an input letter but never moves left. This may model a situation where we have some finite number of probabilistic events and we know that there is a fixed order and number of transitions between them, but with each event being applied an arbitrary number of times. The model is also related to “promise problems” whereby we restrict the decision question to a subset of possible inputs [16]. Letter-bounded languages allow a natural and substantial extension to decision questions on a unary alphabet.

The emptiness and λ -probability problems for PFA on letter-bounded languages were shown to be undecidable for high (finite) dimensional matrices via an encoding of Hilbert’s tenth problem on the solvability of Diophantine equations and Turakainen’s method to transform weighted integer automata to probabilistic automata [25]. These undecidability results also hold for polynomially ambiguous PFA with commutative matrices [4].

The authors of [10] studied decision problems for PFA of various degrees of ambiguity. The degree of ambiguity (finite, polynomial or exponential) of a PFA is a structural property, giving an indication of the number of accepting runs for a given input word. The degree of ambiguity of automata is a well-known and well-studied property in automata theory [26]. The authors of [10] show that the emptiness problem for PFA remains undecidable even for polynomially ambiguous automata (quadratic ambiguity), show **PSPACE**-hardness results for finitely ambiguous PFA and that emptiness is in **NP** for the class of k -ambiguous PFA for every $k > 0$. The emptiness problem for PFA was later shown to be undecidable for linearly ambiguous automata [9].

1.1 Our Contributions

It is natural to consider the decidability of the cutpoint isolation problem for polynomially ambiguous PFA on letter-bounded or commutative languages, given that the (strict) emptiness problems for such automata are undecidable [4]. In the present paper we prove the surprising

result that the cutpoint isolation problem is in fact *decidable*, even if the PFA is exponentially ambiguous, matrices are non-commutative, and the input language is not just the letter-bounded language $a_1^* \cdots a_\ell^*$ but instead a more general letter-bounded context-free language. The results are shown in Table 1.

■ **Table 1** The decidability of problems under different restrictions on the PFA. The main result of this paper is shown in **boldface**. Symbol \implies denotes that decidability is implied by the decidability of the more general model; \impliedby denotes that undecidability is implied by the more restricted model.

Problem	Polynomial ambiguity	Letter-bounded CFL input; Exponential ambiguity	Polynomial ambiguity; letter-bounded input; commutative matrices
(Strict) Emptiness	Undecidable [9, 10, 22] \impliedby	\impliedby	Undecidable [4]
Cutpoint isolation	Undecidable [5, 10]	Decidable	\implies

The result is surprising since in order to solve the cutpoint isolation problem, we must solve two subproblems. Either the cutpoint λ can be reached exactly (the λ -probability problem), or else it can only be approximated arbitrarily closely and is only reached exactly in some limit. As mentioned, the emptiness problem for cutpoint languages is undecidable for polynomially ambiguous PFA on letter-bounded languages, even when all matrices commute [4]. The proof of this result shows a construction of a PFA for which determining if a given $\lambda \in [0, 1]$ is ever reached (i.e., the λ -probability problem) is undecidable. This may at first seem to contradict the results of this paper, since the λ -probability problem is one of the two subproblems to be solved for cutpoint isolation. Why is there no contradiction then? It comes from the fact that as the powers of matrices used in the PFA constructed in [4] increase, the PFA valuation tends towards the limit value λ . Therefore, this λ is always non-isolated and hence the cutpoint isolation problem for such constructed PFA and λ is decidable. However, determining if the PFA ever *exactly* reaches λ is undecidable. So, there is no contradiction with the results of this paper. Our main result is stated as follows.

► **Theorem 1.** *The cutpoint isolation problem for probabilistic finite automata where inputs are constrained to a given letter-bounded context-free language is decidable. Moreover, if the cutpoint is isolated, then a separation bound $\epsilon > 0$ can be computed such that no input word's acceptance probability lies within ϵ of the cutpoint.*

The proof of Theorem 1 is found in Section 3. Our proof technique for showing the decidability of cutpoint isolation for PFA on letter-bounded languages uses the following crucial facts. If a PFA over a letter-bounded context-free language can approach some given cutpoint λ arbitrarily closely, then the PFA can reach λ *exactly* if we allow a subset of the matrices to be taken to one of their “limiting powers”. We use the property that each limiting power (of which there may be finitely many) of a stochastic matrix can be computed in polynomial time (see Lemma 6), as well as a crucial property from linear algebra that dominant eigenvalues (those of strictly largest magnitude) of a stochastic matrix are necessarily of magnitude 1, roots of unity and they have equal geometric and algebraic multiplicities (see Lemma 5). Since the input words of the PFA come from a letter-bounded CFL, we also use the fact that a letter-bounded language is context-free if and only if its Parikh image is a stratified semilinear set (see Proposition 3).

The combination of these ideas allows us to derive Algorithm 1, which works as follows. We initially set all variables as free (rather than fixed), and compute the Parikh image $p(\mathcal{L})$ of the given letter-bounded CFL \mathcal{L} . Using the fact that $p(\mathcal{L})$ is a semilinear set, we compute which letters can be taken to arbitrarily high powers and which letters have fixed finite values.

We then use the technical Proposition 7 which states that if we can reach λ then we can either do so by setting all free variables to an infinite power (which we denote by ω), or else we can compute an integer C such that the value of one of free variables must be less than C . We then either set all free variables as ω in the first case, or nondeterministically choose one of the free variables and assign it a value less than C in the latter case. In the second case we also update the semilinear set and repeat the above procedure until no free variables remain. Finally, we verify that the PFA has exactly the value λ for the chosen values of the variables.

The crucial Proposition 7 is somewhat technical, but relies on splitting a product of stochastic matrices into a summation involving dominant and subdominant eigenvalues (a subdominant eigenvalue being one with magnitude strictly less than 1) and then applying the spectral decomposition or Jordan normal form of each stochastic matrix in order to derive the constant C which bounds the value of one of free variables.

Combining our proof technique with a result of Rabin [23], we derive the following result.

► **Corollary 9.** *The emptiness problem is decidable for probabilistic finite automata on letter-bounded context-free languages when the cutpoint is isolated.*

The undecidability of the emptiness problem for PFA over letter-bounded inputs shown in [4] therefore only applies when the cutpoint is non-isolated.

The provided algorithm is nondeterministic in nature although we do not have an upper bound on its complexity. We can however provide the following lower bound via an adaptation of a proof technique from [4] which proved the NP-hardness of the injectivity problem for linearly ambiguous three-state probabilistic finite automata over letter-bounded languages.

► **Theorem 12.** *Cutpoint isolation is NP-hard for 3-state PFA on letter-bounded inputs.*

Our procedure also allows us to answer some equivalent problems (in Section 5), for example: given a PFA, $\lambda \in [0, 1]$ and a maximum number $k \in \mathbb{N}$ of alternations between input letters, determine if λ is isolated. We also prove the value-1 problem is decidable over letter-bounded context-free language inputs.

2 Preliminaries

2.1 Probabilistic Finite Automata on Letter-Bounded Inputs

We denote by $\mathbb{F}^{n \times n}$ the set of all $n \times n$ matrices over some field \mathbb{F} . We will primarily be interested in rational matrices. We use a nonstandard form of *Dirac bra-ket notation* in several calculations, to simplify the notation in some complex formulae. If $u = (u_1, \dots, u_n)^\top \in \mathbb{C}^n$ is a column vector, then we write $|u\rangle = u$ and $\langle u| = u^\top$ where u^\top denotes the transpose of u , i.e., $\langle u| = (u_1, \dots, u_n)$. Note that Dirac bra-ket notation ordinarily defines that $\langle u| = u^*$ where u^* denotes the *conjugate* transpose of u , however we will not use this notion at any point. Note that $|u\rangle\langle v|$ is just a rank 1 matrix $u^\top v$. We use $\langle e_i|$ and $|e_i\rangle$ to denote the i 'th basis row/column vector respectively.

A PFA \mathcal{P} with n states over an alphabet Σ is defined as $\mathcal{A} = (\langle u|, \{M_a | a \in \Sigma\}, |v\rangle)$ where $\langle u| \in \mathbb{R}^n$ is the initial probability distribution; $|v\rangle \in \{0, 1\}^n$ is the final state vector and each $M_a \in \mathbb{R}^{n \times n}$ is a (row) stochastic matrix. For a word $w = w_1 w_2 \cdots w_k \in \Sigma^*$, we define the acceptance probability $f_{\mathcal{P}} : \Sigma^* \rightarrow \mathbb{R}$ of \mathcal{P} as:

$$f_{\mathcal{P}}(w) = \langle u| M_{w_1} M_{w_2} \cdots M_{w_k} |v\rangle,$$

which denotes the acceptance probability of w .¹

¹ Some authors interchange the order of u and v and use column stochastic matrices, although the two definitions are trivially isomorphic.

For any $\lambda \in [0, 1]$ and PFA \mathcal{P} over alphabet Σ , we define a cutpoint language to be: $L_{\geq \lambda}(\mathcal{P}) = \{w \in \Sigma^* \mid f_{\mathcal{P}}(w) \geq \lambda\}$, and a strict cutpoint language $L_{> \lambda}(\mathcal{P})$ by replacing \geq with $>$. The (strict) emptiness problem for a cutpoint language is to determine if $L_{\geq \lambda}(\mathcal{P}) = \emptyset$ (resp. $L_{> \lambda}(\mathcal{P}) = \emptyset$). Our main focus is on the *cutpoint isolation problem*, now defined.

► **Problem 2** (Cutpoint isolation). *Given a PFA \mathcal{P} and cutpoint $\lambda \in [0, 1]$, determine if for each $\epsilon > 0$ there exists some $w \in \Sigma^*$ such that $|f_{\mathcal{P}}(w) - \lambda| < \epsilon$.*

Let $\Sigma = \{a_1, a_2, \dots, a_\ell\}$ be an alphabet with $\ell > 0$ distinct letters. A language \mathcal{L} is called *letter-bounded* if $\mathcal{L} \subseteq a_1^* a_2^* \dots a_\ell^*$. If \mathcal{L} is letter-bounded and also a context-free language, then it is called a letter-bounded context-free language. We are interested in cutpoint isolation for PFA whose inputs come from a given letter-bounded context-free language.

For a letter-bounded language $\mathcal{L} \subseteq a_1^* a_2^* \dots a_\ell^*$, define its *Parikh image*² as

$$p(\mathcal{L}) = \{(k_1, \dots, k_\ell) : a_1^{k_1} a_2^{k_2} \dots a_\ell^{k_\ell} \in \mathcal{L}\}.$$

Recall that a subset $Q \subseteq \mathbb{N}^\ell$ is called *linear* if there are vectors $q_0, q_1, \dots, q_r \in \mathbb{N}^\ell$ such that

$$Q = \{q_0 + t_1 q_1 + \dots + t_r q_r : t_1, \dots, t_r \in \mathbb{N}\}.$$

We say that a linear set Q is *stratified* if for each $i \geq 1$ the vector q_i has at most two nonzero coordinates, and for any $i, j \geq 1$ if both q_i and q_j have two nonzero coordinates, $i_1 < i_2$ and $j_1 < j_2$, respectively, then their order is *not* $i_1 < j_1 < i_2 < j_2$, i.e., they are not interlaced. A finite union of linear sets is called a *semilinear set*, and a finite union of stratified linear sets is called a *stratified semilinear set*.

We will need the following classical fact about context-free languages.

► **Proposition 3.** *If \mathcal{L} is a context-free language, then its Parikh image $p(\mathcal{L})$ is a semilinear set that can be effectively constructed from the definition of \mathcal{L} [21].*

► **Remark 4.** There is a nice characterization of the letter-bounded context-free languages. Namely, a letter-bounded language $\mathcal{L} \subseteq a_1^* a_2^* \dots a_\ell^*$ is context-free if and only if $p(\mathcal{L})$ is a stratified semilinear set [14, 15].

Let $A_1, \dots, A_\ell \in \mathbb{Q}^{n \times n}$ be row stochastic matrices. Let $u \in \mathbb{Q}^n$ be a stochastic vector (the initial vector) and $v \in \{0, 1\}^n$ (the final state vector). Let $\mathcal{L} \subseteq a_1^* a_2^* \dots a_\ell^*$ be a letter-bounded context-free language, and let $\lambda \in [0, 1]$ be a cutpoint for which we want to decide if it is isolated or not, that is, whether λ belongs to the closure of $\{\langle u \mid A_1^{k_1} A_2^{k_2} \dots A_\ell^{k_\ell} \mid v \rangle : a_1^{k_1} a_2^{k_2} \dots a_\ell^{k_\ell} \in \mathcal{L}\}$.

If λ is not isolated, then there are two scenarios: either there exists $k_1, k_2, \dots, k_\ell \in \mathbb{N}$ such that $\langle u \mid A_1^{k_1} A_2^{k_2} \dots A_\ell^{k_\ell} \mid v \rangle = \lambda$, or else λ is never reached but only approached arbitrarily closely. In the second case there is a sequence of tuples $\{(k_1^m, k_2^m, \dots, k_\ell^m)\}_{m=1}^\infty$ such that

$$\lambda = \lim_{m \rightarrow \infty} \langle u \mid A_1^{k_1^m} A_2^{k_2^m} \dots A_\ell^{k_\ell^m} \mid v \rangle$$

and, furthermore, for every $t \in \{1, \dots, \ell\}$, either $k_t^m = k_t^1$ for all $m \geq 1$, i.e. k_t^m is fixed, or k_t^m is strictly increasing and $A_t^{k_t^m}$ converges to a limit as $m \rightarrow \infty$.

² In general, the Parikh image of $\mathcal{L} \subseteq \Sigma^*$ is defined as $p(\mathcal{L}) = \{(|w|_{a_1}, \dots, |w|_{a_\ell}) : w \in \mathcal{L}\}$ where $|w|_{a_i}$ denotes the number of occurrences of letter a_i in word w .

We will use the notation A^ω to denote the set of all limits of the sequence $\{A^k\}_{k=1}^\infty$ (see Lemma 6 below for a detailed explanation). It follows that if λ is not isolated, then there exists a choice of variables $k_1, k_2, \dots, k_\ell \in \mathbb{N} \cup \{\omega\}$ such that

$$\lambda \in \langle u | A_1^{k_1} A_2^{k_2} \cdots A_\ell^{k_\ell} | v \rangle.$$

Note that if $k_t = \omega$, then A_t^ω is a finite set. In this case we substitute all limits of A_t^ω in the above formula, and so $\langle u | A_1^{k_1} A_2^{k_2} \cdots A_\ell^{k_\ell} | v \rangle$ also becomes a finite set.³

2.2 Algebraic numbers

A complex number α is *algebraic* if it is a root of a polynomial $p \in \mathbb{Z}[x]$. The *defining polynomial* $p_\alpha \in \mathbb{Z}[x]$ for α is the unique polynomial of least degree with positive leading coefficient such that the coefficients of p_α do not have a common factor and $p_\alpha(\alpha) = 0$. The *degree* and *height* of α are defined to be that of p_α .

In order to do computations with algebraic numbers we use their standard representations. Namely, an algebraic number can be represented by its defining polynomial and a sufficiently good complex rational approximation. More precisely, α will be represented by a tuple (p_α, a, b, r) , where $p_\alpha \in \mathbb{Z}[x]$ is the defining polynomial for α and $a, b, r \in \mathbb{Q}$ are such that α is the unique root of p_α inside the circle in \mathbb{C} with centre $a + bi$ and radius r . As shown in [19], if $\alpha \neq \beta$ are roots of $p \in \mathbb{Z}[x]$, then $|\alpha - \beta| > \frac{\sqrt{6}}{d^{(d+1)/2} H^{d-1}}$, where d and H are the degree and height of p , respectively. So, if we require r to be smaller than half of this bound, the above representation is well-defined.

Let $\|\alpha\|$ be the size of the standard representation of α , that is, the total bit size of a, b, r and the coefficients of p_α . It is well-known fact that for given algebraic numbers α and β , one can compute $1/\alpha$, $\bar{\alpha}$ and $|\alpha|$ in time polynomial in $\|\alpha\|$, and one can compute $\alpha + \beta$ and $\alpha\beta$ and decide whether $\alpha = \beta$ in time polynomial in $\|\alpha\| + \|\beta\|$. Moreover, for a *real* algebraic α , deciding whether $\alpha > 0$ can be done in time polynomial in $\|\alpha\|$. Finally, there is a polynomial time algorithm that for a given $p \in \mathbb{Z}[x]$ computes the standard representations of all roots of p . For more information on efficient algorithmic computations with algebraic numbers the reader is referred to [1, 8, 17, 20].

2.3 Spectral decomposition and Jordan normal forms

We define the spectrum (set of eigenvalues) of $A \in \mathbb{R}^{n \times n}$ as $\sigma(A) = \{\lambda_1, \dots, \lambda_n\}$ arranged in monotonically nonincreasing order, i.e. $|\lambda_i| \geq |\lambda_j|$ for all $1 \leq i < j \leq n$ and we define $\hat{\sigma}(A) \subseteq \sigma(A)$ as the set of eigenvalues of A of absolute value 1. We call eigenvalues $\hat{\sigma}(A)$ *dominant eigenvalues* and eigenvalues $\sigma(A) \setminus \hat{\sigma}(A)$ *subdominant eigenvalues*.

Given $A = (a_{ij}) \in \mathbb{F}^{m \times m}$ and $B \in \mathbb{F}^{n \times n}$, we define the direct sum $A \oplus B$ of A and B by:

$$A \oplus B = \left[\begin{array}{c|c} A & \mathbf{0}_{m,n} \\ \hline \mathbf{0}_{n,m} & B \end{array} \right], \text{ where } \mathbf{0}_{n,m} \text{ is the } n \times m \text{ zero matrix.}$$

We will use both the *spectral decomposition theorem* and the *Jordan normal form* of stochastic matrices in later proofs. For background, see [11].

³ If $\{A^k\}_{k=1}^\infty$ has a unique limit A' , then we will identify the set $A^\omega = \{A'\}$ with the matrix A' and write $A^\omega = A'$. Also, if all k_t 's are finite or if all limits are unique, we identify the number $\langle u | A_1^{k_1} A_2^{k_2} \cdots A_\ell^{k_\ell} | v \rangle$ with the one element set $\{\langle u | A_1^{k_1} A_2^{k_2} \cdots A_\ell^{k_\ell} | v \rangle\}$.

Let $A_i \in \mathbb{Q}^{n \times n}$ be a matrix (we use notation A_i since it will prove useful in the proof of Proposition 7), and let $\{\lambda_{i,1}, \dots, \lambda_{i,n_i}\}$ be the eigenvalues of A_i listed according to their *geometric* multiplicities⁴. Then A_i can be written in *Jordan normal form* $A_i = S_i^{-1}(J_{\ell_{i,1}}(\lambda_{i,1}) \oplus \dots \oplus J_{\ell_{i,n_i}}(\lambda_{i,n_i}))S_i$, where S_i is an invertible matrix ($\det(S_i) \neq 0$) and $J_{\ell_{i,j}}(\lambda_{i,j})$ is a $\ell_{i,j} \times \ell_{i,j}$ *Jordan block* for $1 \leq j \leq n_i \leq n$, with n_i the number of Jordan blocks of A_i and $\ell_{i,j}$ the size of the Jordan block corresponding to eigenvalue $\lambda_{i,j}$, such that $\ell_{i,1} + \dots + \ell_{i,n_i} = n$. Jordan block $J_{\ell_{i,j}}(\lambda_{i,j})$ corresponds to the j^{th} eigenvalue $\lambda_{i,j}$ of A_i and has the form:

$$J_{\ell_{i,j}}(\lambda_{i,j}) = \begin{pmatrix} \lambda_{i,j} & 1 & 0 & \cdots & 0 \\ 0 & \lambda_{i,j} & 1 & \cdots & 0 \\ 0 & 0 & \lambda_{i,j} & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & \lambda_{i,j} \end{pmatrix} \in \mathbb{C}^{\ell_{i,j} \times \ell_{i,j}}$$

The matrix S_i contains the *generalised* eigenvectors of A_i . Noting that $\binom{x}{y} = 0$ if $y > x$, we now see that

$$\begin{aligned} J_{\ell_{i,j}}(\lambda_{i,j})^{k_i} &= \begin{pmatrix} \lambda_{i,j}^{k_i} & \binom{k_i}{1} \lambda_{i,j}^{k_i-1} & \binom{k_i}{2} \lambda_{i,j}^{k_i-2} & \cdots & \binom{k_i}{\ell_{i,j}-1} \lambda_{i,j}^{k_i-(\ell_{i,j}-1)} \\ 0 & \lambda_{i,j}^{k_i} & \binom{k_i}{1} \lambda_{i,j}^{k_i-1} & \cdots & \binom{k_i}{\ell_{i,j}-2} \lambda_{i,j}^{k_i-(\ell_{i,j}-2)} \\ 0 & 0 & \lambda_{i,j}^{k_i} & \cdots & \binom{k_i}{\ell_{i,j}-3} \lambda_{i,j}^{k_i-(\ell_{i,j}-3)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & \lambda_{i,j}^{k_i} \end{pmatrix} \in \mathbb{C}^{\ell_{i,j} \times \ell_{i,j}} \\ &= \sum_{0 \leq m \leq \ell_{i,j}-1} \lambda_{i,j}^{k_i-m} \binom{k_i}{m} \left(\sum_{1 \leq p \leq \ell_{i,j}-m} |e_p\rangle \langle e_{m+p}| \right) \end{aligned} \quad (1)$$

The *spectral decomposition* of a matrix is a special case of the Jordan normal form. Namely, any *diagonalizable* matrix $A_i \in \mathbb{Q}^{n \times n}$ can be written as

$$A_i = S_i^{-1}(\lambda_{i,1} \oplus \dots \oplus \lambda_{i,n})S_i = \sum_{j=1}^n \lambda_{i,j} |v_{i,j}\rangle \langle u_{i,j}|, \quad (2)$$

where $\sigma(M) = \{\lambda_{i,1}, \dots, \lambda_{i,n}\}$ is the set of eigenvalues of A_i , $|v_{i,j}\rangle$ is the j^{th} column of S_i^{-1} and $\langle u_{i,j}|$ is the j^{th} row of S_i . Thus we have $A_i^k = \sum_{j=1}^n \lambda_{i,j}^k |v_{i,j}\rangle \langle u_{i,j}|$.

We will also require the following technical lemma concerning the dominant eigenvalues of stochastic matrices.

► **Lemma 5** ([11, Theorem 6.5.3]). *Let λ be a dominant eigenvalue of a stochastic matrix $A \in \mathbb{R}^{n \times n}$. Then λ is a root of unity of order no more than n . Moreover, the geometric multiplicity of λ is equal to its algebraic multiplicity. In other words, the Jordan blocks that correspond to λ have size 1×1 .*

We also require the following lemma.

⁴ Note that n_i is the number of linearly independent eigenvectors of A_i or the number of Jordan blocks in the Jordan normal form of A_i . The matrix A_i is diagonalizable if and only if $n_i = n$. Jordan normal forms are unique up to permutations of the Jordan blocks.

► **Lemma 6.** *For any stochastic matrix A , the sequence $\{A^k\}_{k=1}^\infty$ has a finite number of limits. Namely, there exist a computable constant d such that, for each $r = 0, \dots, d-1$, the subsequence $\{A^{dm+r}\}_{m=1}^\infty$ converges to a limit, and this limit can be computed in polynomial time given d and r .*

Proof. Let A be a stochastic matrix. As shown in [7], we can compute in polynomial time the Jordan normal form of A and a transformation matrix S such that $A = S^{-1}JS$. Note that A may have complex eigenvalues, so all computations are done using standard representations of algebraic numbers as explained in Section 2.2.

By Lemma 5, all dominant eigenvalues of A are roots of unity of orders no more than n , and their Jordan blocks have size 1×1 . If λ is a root of unity of order p , then $\{\lambda^k\}_{k=1}^\infty$ is a periodic sequence with period p . On the other hand, if $J_\ell(\lambda)$ is a Jordan block corresponding to an eigenvalue λ such that $|\lambda| < 1$, then $\lim_{k \rightarrow \infty} J_\ell(\lambda)^k$ is equal to the zero matrix.

Let d be the least common multiple of the orders of the roots of unity among the eigenvalues of A . Now if λ is a dominant eigenvalue of A , then the values of $\lambda^{dm+r} = \lambda^r$ do not depend on m , where $r = 0, \dots, d-1$. Hence J^{dm+r} converges to a limit when $m \rightarrow \infty$. This limit is equal to a matrix J' obtained from J by replacing all dominant λ with λ^r and all Jordan blocks corresponding to subdominant eigenvalues with zero matrices. So, $\lim_{m \rightarrow \infty} A^{dm+r} = S^{-1}J'S$.

This shows that $\{A^k\}_{k=1}^\infty$ has at most d limits. Finally, we note that d may be exponential in the dimension of A . However, if $\{A^k\}_{k=1}^\infty$ has a single limit, then this limit can be computed in polynomial time. ◀

3 Decidability of Cutpoint Isolation

In this section we will give a proof of Theorem 1 which is our main result. The crucial ingredient of our proof is the following technical proposition which will be proven in Section 4.

► **Proposition 7.** *Let $J = \{1, 2, \dots, \ell\}$ be indices, $\lambda \in [0, 1]$ a cutpoint, and let $J_F \subseteq J$ be such that k_t is a free variable, for $t \in J_F$, and k_t is assigned a fixed finite value, for $t \in J \setminus J_F$. Then*

- *either $\lambda \in \langle u | A_1^{k_1} A_2^{k_2} \dots A_\ell^{k_\ell} | v \rangle$, where $k_t = \omega$ for all $t \in J_F$,*
- *or else there exists a constant $C > 0$ such that $\lambda \in \langle u | A_1^{k_1} A_2^{k_2} \dots A_\ell^{k_\ell} | v \rangle$ implies $k_t < C$ for at least one $t \in J_F$.*

Moreover, we can decide whether the first case holds and compute the constant C in the second case.

Below we give a high-level description of the main algorithm (Algorithm 1), which gives a formal proof of Theorem 1, and explain how Proposition 7 is used there.

Let $\mathcal{L} \subseteq a_1^* a_2^* \dots a_\ell^*$ be a given letter-bounded CFL. We start by considering all indices $J = \{1, \dots, \ell\}$ as *free* (i.e. their value is not fixed and they will later be given a fixed value from $\mathbb{N} \cup \{\omega\}$) and iteratively fix them until no free indices remain. We first use Parikh and Ginsburg's results (Proposition 3) to compute the Parikh image $p(\mathcal{L})$. Then we nondeterministically choose a linear subset Q and use it to determine the indices which can be taken to arbitrary high values while staying within Q . These indices will correspond to the “free variables” in the algorithm.

Let J_F be a set of such indices (which will be called R in Algorithm 1). We then set k_t for $t \in J \setminus J_F$ to appropriate finite values, while k_t with $t \in J_F$ remain free variables. We wish to determine if there is a choice of $k_t \in \mathbb{N} \cup \{\omega\}$ for $t \in J_F$ such that

$$\lambda \in \langle u | A_1^{k_1} A_2^{k_2} \dots A_\ell^{k_\ell} | v \rangle,$$

that is, whether λ can be reached by setting each free variable k_t either to some finite value or else to ω , an “infinite” power. Proposition 7 then tells us that either all free variables should be set at ω in order to reach λ (and this is decidable), or else there exists a computable constant C such that *if* we can reach λ by some choice of these free variables, then some $k_t < C$ for an index $t \in J_F$.

In the first case, we set all free variables to ω . In the second case, we nondeterministically choose some free variable, fix its value in the range $[0, C)$ and then update our linear set Q to satisfy a new constraint. The procedure repeats iteratively until all free variables have been assigned a fixed value. The algorithm then verifies if this choice of variables gives a solution.

■ **Algorithm 1** Nondeterministic algorithm deciding whether a given cutpoint is isolated.

Stage one (Nondeterministic iterative fixing of free variables):

Let $T = J = \{1, 2, \dots, \ell\}$.

Compute the Parikh image $p(\mathcal{L})$ and nondeterministically choose one of its finitely many linear subsets $Q = \{q_0 + t_1q_1 + \dots + t_rq_r : t_1, \dots, t_r \in \mathbb{N}\} \subseteq p(\mathcal{L})$.⁵

while $T \neq \emptyset$ **do**

Let R be the set of indices $j \in T$ such that at least one q_i with $i \geq 1$ has a nonzero j th coordinate.⁶

For each $j \in T \setminus R$, the j th coordinate of all vectors from Q is equal to the j th coordinate of q_0 .⁷ So, we set k_j for $j \in T \setminus R$ to be the j th coordinate of q_0 .

Then, for $j \in R$, compute the limits A_j^ω with indices respecting set Q (see Remark 8 below for details).

Check whether $\lambda \in \langle u | A_1^{k_1} A_2^{k_2} \dots A_\ell^{k_\ell} | v \rangle$, where $k_j = \omega$ for all $j \in R$.

If yes, return True and stop.

Otherwise, assuming all indices k_j for $j \in J \setminus R$ are fixed and R is the set of free variables, use Proposition 7 to compute the constant $C > 0$ such that

$\lambda \in \langle u | A_1^{k_1} A_2^{k_2} \dots A_\ell^{k_\ell} | v \rangle$ implies $k_j < C$ for at least one $j \in R$.

Then nondeterministically choose $j \in R$, fix $k_j \in [0, C)$ and set $T \leftarrow R \setminus \{j\}$.

Next, for the chosen index j , find those indices i in $\{1, \dots, r\}$ for which q_i has a nonzero j th coordinate. Without loss of generality, suppose $\{1, \dots, s\}$ are these indices.

Fixing $k_j \in [0, C)$, restricts the parameters t_1, \dots, t_s in Q to a finite set of possible values since the vector $q_0 + t_1q_1 + \dots + t_sq_s$ must have k_j in its j th coordinate.

Nondeterministically choose one of these values for t_1, \dots, t_s or return False and stop, if such a choice is impossible.

Let $Q \leftarrow \{(q_0 + t_1q_1 + \dots + t_sq_s) + t_{s+1}q_{s+1} + \dots + t_rq_r : t_{s+1}, \dots, t_r \in \mathbb{N}\}$.⁸

Stage two (Verifying the computation):

At this stage we have fixed all variables k_1, k_2, \dots, k_ℓ to some finite values.

Compute $\langle u | A_1^{k_1} A_2^{k_2} \dots A_\ell^{k_\ell} | v \rangle$ for the obtained values of $k_1, \dots, k_\ell \in \mathbb{N}$.

Return True if $\lambda \in \langle u | A_1^{k_1} A_2^{k_2} \dots A_\ell^{k_\ell} | v \rangle$ or False, otherwise.

End.

► **Remark 8.** To compute the limits A_j^ω with indices respecting set Q , note that the projection of Q on the j th coordinate is equal to $\{q_{0,j} + t_1q_{1,j} + \dots + t_rq_{r,j} : t_1, \dots, t_r \in \mathbb{N}\}$, where $q_{i,j}$ is the j th coordinate of q_i . The set $\langle q_{1,j}, \dots, q_{r,j} \rangle = \{t_1q_{1,j} + \dots + t_rq_{r,j} : t_1, \dots, t_r \in \mathbb{N}\}$ is a finitely generated subsemigroups of $(\mathbb{N}, +)$. Let $d = \gcd(q_{1,j}, \dots, q_{r,j})$, then there is a number $s > 0$ such that for any $t \geq s$, we have $t \in \langle q_{1,j}, \dots, q_{r,j} \rangle$ if and only if d divides t .

22:10 Decidability of Cutpoint Isolation for PFA on Letter-Bounded Inputs

This is a well-known property of the subsemigroups of $(\mathbb{N}, +)$ [24]. Thus the limits A_j^ω with indices respecting Q are equal to $A_j^\omega = A_j^{q_0, j} (A_j^d)^\omega$, where the limits $(A_j^d)^\omega$ are computed using Lemma 6.

It remains to prove that we can compute a separation bound $\epsilon > 0$ between λ and the closest acceptance probability of \mathcal{P} for any input word $w \in \mathcal{L}$. Algorithm 1 has two stopping conditions, either by returning True in **Stage one** (which we discount since it implies the cutpoint is not isolated), or else after **Stage two**.

Algorithm 1 has two sources of nondeterminism in **Stage one**: in the choice of linear subset Q and then during the while loop in the choice of $j \in R$ and $k_j \in [0, C)$. We will evaluate every choice of linear subset Q and every choice of j and k_j to cover all possible cases, updating a global variable ϵ at the end of every nondeterministic branch. Initially, we set $\epsilon \leftarrow \infty$, and let ϵ_1, ϵ_2 be additional global variables that are set $\epsilon_1 \leftarrow \epsilon_2 \leftarrow \infty$ at the beginning of every nondeterministic branch.

During the execution of **Stage one** we use Proposition 7 to compute C such that if all free variables are above C then we are at least some $\epsilon' > 0$ away from λ . Note that ϵ' is less than half of the distance between λ and some limit values. For each iteration of the while loop, we set $\epsilon_1 \leftarrow \min\{\epsilon_1, \epsilon'\}$ to keep track of the minimal value. During **Stage two**, all variables have a fixed finite value, and we set $\epsilon_2 \leftarrow |\langle u | A_1^{k_1} A_2^{k_2} \cdots A_\ell^{k_\ell} | v \rangle - \lambda|$ which is greater than zero assuming λ is isolated. Finally, we set $\epsilon \leftarrow \min\{\epsilon, \epsilon_1, \epsilon_2\} > 0$.

After inspecting all possible nondeterministic runs of the algorithm, the obtained value of ϵ gives us the separation bound. Indeed, during the execution of the above procedure, ϵ is updated to the minimum of ϵ_1 and ϵ_2 , where ϵ_1 is less than half of the distance between λ and some limit values and ϵ_2 keeps track of the distance between λ and the values $\langle u | A_1^{k_1} A_2^{k_2} \cdots A_\ell^{k_\ell} | v \rangle$ when each index k_j is less than the corresponding constant C .

4 Proof of Proposition 7

We begin with a proof sketch. Since each A_i is stochastic, $\hat{\sigma}(A_i)$ contains at least one eigenvalue 1 and all other eigenvalues in $\hat{\sigma}(A_i)$ are roots of unity by Lemma 5. All eigenvalues in $\sigma(A_i) \setminus \hat{\sigma}(A_i)$ have absolute value strictly smaller than 1. Our approach is to rewrite the expression

$$\langle u | A_1^{k_1} A_2^{k_2} \cdots A_\ell^{k_\ell} | v \rangle \quad (3)$$

into the sum of two terms (which will be denoted by S_0 and S_1) such that S_0 determines the limit behaviour as all free variables tend towards infinity, since they control only dominant eigenvalues, while S_1 is vanishing, since at least one free variable controls a subdominant eigenvalue. We can then reason that if all free variables simultaneously become larger, then Eqn (3) tends towards a set of computable limits with some vanishing terms. Therefore we can determine either that we can reach λ when all free variables are ω , or else we can prove that Eqn (3) is within any $\epsilon > 0$ of a limit value once all free variables are sufficiently large, which proves the proposition (by setting ϵ as less than the smallest difference from a limit value and λ). We now proceed with the formal details.

First, we consider the simpler case when all matrices are diagonalizable and then show how to extend this argument to the general case.

Diagonalizable matrices. Let us first assume that all matrices are diagonalizable. By the spectral decomposition theorem (see Eqn (2)), we may write a matrix $A_i^{k_i}$ as:

$$A_i^{k_i} = \sum_{j=1}^n \lambda_{i,j}^{k_i} |v_{i,j}\rangle \langle u_{i,j}|, \quad (4)$$

where $\{\lambda_{i,1}, \dots, \lambda_{i,n}\}$ are the eigenvalues of A_i repeated according to their multiplicities, and the vectors $|v_{i,j}\rangle$ and $\langle u_{i,j}|$, for $1 \leq j \leq n$, are related to the eigenvectors of A_i . Now, we can write:

$$\begin{aligned} \langle u| A_1^{k_1} A_2^{k_2} \cdots A_\ell^{k_\ell} |v\rangle &= \langle u| \left(\prod_{i=1}^{\ell} \left(\sum_{j=1}^n \lambda_{i,j}^{k_i} |v_{i,j}\rangle \langle u_{i,j}| \right) \right) |v\rangle \\ &= \sum_{j_1, \dots, j_\ell \in [1, n]} \lambda_{1,j_1}^{k_1} \lambda_{2,j_2}^{k_2} \cdots \lambda_{\ell,j_\ell}^{k_\ell} \langle u|v_{1,j_1}\rangle \langle u_{1,j_1}|v_{2,j_2}\rangle \langle u_{2,j_2}| \cdots |v_{\ell,j_\ell}\rangle \langle u_{\ell,j_\ell}|v\rangle \end{aligned}$$

Let us thus define $\Theta_{j_1, \dots, j_\ell} = \langle u|v_{1,j_1}\rangle \langle u_{1,j_1}|v_{2,j_2}\rangle \langle u_{2,j_2}| \cdots |v_{\ell,j_\ell}\rangle \langle u_{\ell,j_\ell}|v\rangle$. The above sum can be split in two: the first summand containing terms where only dominant eigenvalues are to the power of free variables, and the second containing terms with at least one subdominant eigenvalue to the power of a free variable (these two terms are labelled S_0 and S_1 below). This is a useful decomposition since any term which contains a subdominant eigenvalue taken to the power of a free variable will tend towards zero as the values of all free variables (simultaneously) increase. Thus we can write $\langle u| A_1^{k_1} A_2^{k_2} \cdots A_\ell^{k_\ell} |v\rangle = S_0 + S_1$, where

$$\begin{aligned} S_0 &= \sum_{\substack{j_1, \dots, j_\ell \in [1, n] \\ \forall t \in J_F : |\lambda_{t,j_t}| = 1}} \lambda_{1,j_1}^{k_1} \lambda_{2,j_2}^{k_2} \cdots \lambda_{\ell,j_\ell}^{k_\ell} \Theta_{j_1, \dots, j_\ell}, \\ S_1 &= \sum_{\substack{j_1, \dots, j_\ell \in [1, n] \\ \exists t \in J_F : |\lambda_{t,j_t}| < 1}} \lambda_{1,j_1}^{k_1} \lambda_{2,j_2}^{k_2} \cdots \lambda_{\ell,j_\ell}^{k_\ell} \Theta_{j_1, \dots, j_\ell}. \end{aligned}$$

By Lemma 5 the dominant eigenvalues are roots of unity, and so S_0 assumes only finitely many different values as k_t with $t \in J_F$ vary, while k_t with $t \in J \setminus J_F$ are fixed.

Suppose S_1 is not an empty sum since otherwise $S_1 = 0$. Then there exists $t \in J_F$ and $j_t \in [1, n]$ such that $|\lambda_{t,j_t}| < 1$. Let ρ be the maximum among such values, that is,

$$\rho = \max\{|\lambda_{t,j}| : t \in J_F, j \in [1, n] \text{ and } |\lambda_{t,j}| < 1\}.$$

Suppose $k_t \geq C$ for $t \in J_F$, where C is some constant to be chosen later. Then S_1 can be estimated as follows: since for every choice of j_1, \dots, j_ℓ in the summation S_1 there is $t \in J_F$ with $|\lambda_{t,j_t}| \leq \rho < 1$ and $|\lambda_{i,j}| \leq 1$ for all other $\lambda_{i,j}$, we have

$$|S_1| \leq C_1 \rho^C, \quad \text{where } C_1 = \sum_{\substack{j_1, \dots, j_\ell \in [1, n] \\ \exists t \in J_F : |\lambda_{t,j_t}| < 1}} |\Theta_{j_1, \dots, j_\ell}|.$$

Notice that for any rational $\epsilon > 0$, we can compute $C \in \mathbb{N}$ such that $|S_1| \leq C_1 \rho^C < \epsilon$. Now, S_0 gives a finite number of limit values for $\langle u| A_1^{k_1} A_2^{k_2} \cdots A_\ell^{k_\ell} |v\rangle$. If λ is not equal to any of them, then choose $\epsilon > 0$ to be less than half the minimal distance between λ and those limit values. Using this ϵ , we compute C as above. By definition of C , if all $k_t \geq C$ for $t \in J_F$, then the distance between $\langle u| A_1^{k_1} A_2^{k_2} \cdots A_\ell^{k_\ell} |v\rangle$ and one of the limit values of S_0 is less than ϵ . Thus $\langle u| A_1^{k_1} A_2^{k_2} \cdots A_\ell^{k_\ell} |v\rangle$ cannot be equal to λ when all $k_t \geq C$ for $t \in J_F$. Hence if $\lambda = \langle u| A_1^{k_1} A_2^{k_2} \cdots A_\ell^{k_\ell} |v\rangle$, then there is $t \in J_F$ such that $k_t < C$.

The general case. We now show how to extend the proof to the case when some matrices are non-diagonalizable.

Let $a_{i,j} = \sum_{s=1}^{j-1} \ell_{i,s}$ be the sum of the sizes of the first $j-1$ Jordan blocks of matrix A_i , so that $a_{i,1} = 0, a_{i,2} = \ell_{i,1}, a_{i,3} = \ell_{i,1} + \ell_{i,2}$ etc. Then we see that by using Eqn (1), $A_i^{k_i} = S_i^{-1} (J_{\ell_{i,1}}(\lambda_{i,1})^{k_i} \oplus \cdots \oplus J_{\ell_{i,n_i}}(\lambda_{i,n_i})^{k_i}) S_i$ has the form

$$\begin{aligned}
 & S_i^{-1} \left(\sum_{1 \leq j \leq n_i} \sum_{0 \leq m \leq \ell_{i,j}-1} \lambda_{i,j}^{k_i-m} \binom{k_i}{m} \left(\sum_{1 \leq p \leq \ell_{i,j}-m} |e_{a_{i,j}+p}\rangle \langle e_{a_{i,j}+m+p}| \right) \right) S_i \\
 &= \sum_{1 \leq j \leq n_i} \sum_{0 \leq m \leq \ell_{i,j}-1} \lambda_{i,j}^{k_i-m} \binom{k_i}{m} \left(\sum_{1 \leq p \leq \ell_{i,j}-m} S_i^{-1} |e_{a_{i,j}+p}\rangle \langle e_{a_{i,j}+m+p}| S_i \right) \\
 &= \sum_{1 \leq j \leq n_i} \sum_{0 \leq m \leq \ell_{i,j}-1} \lambda_{i,j}^{k_i-m} \binom{k_i}{m} \left(\sum_{1 \leq p \leq \ell_{i,j}-m} |v_{i,a_{i,j}+p}\rangle \langle u_{i,a_{i,j}+m+p}| \right),
 \end{aligned}$$

where $S_i = \sum_{q=1}^n |e_q\rangle \langle u_{i,q}|$ and $S_i^{-1} = \sum_{q=1}^n |v_{i,q}\rangle \langle e_q|$, with e_q the q^{th} basis vector. Here we used the property that $\langle e_i | e_j \rangle = 0$ for any $i \neq j$. We may now compute that:

$$\begin{aligned}
 & \langle u | A_1^{k_1} A_2^{k_2} \cdots A_\ell^{k_\ell} | v \rangle \\
 &= \langle u | \prod_{i=1}^{\ell} \left(\sum_{1 \leq j \leq n_i} \sum_{0 \leq m \leq \ell_{i,j}-1} \lambda_{i,j}^{k_i-m} \binom{k_i}{m} \left(\sum_{1 \leq p \leq \ell_{i,j}-m} |v_{i,a_{i,j}+p}\rangle \langle u_{i,a_{i,j}+m+p}| \right) \right) | v \rangle \\
 &= \langle u | \prod_{i=1}^{\ell} \left(\sum_{1 \leq j \leq n_i} \sum_{0 \leq m \leq \ell_{i,j}-1} \lambda_{i,j}^{k_i-m} \binom{k_i}{m} \Psi_{i,j,m} \right) | v \rangle \\
 &= \sum_{\substack{j_1, \dots, j_\ell \mid j_q \in [1, n_q] \\ m_1, \dots, m_\ell \mid m_q \in [0, \ell_{q,j_q}-1]}} \left(\prod_{1 \leq t \leq \ell} \lambda_{t,j_t}^{k_t-m_t} \binom{k_t}{m_t} \right) \langle u | \Psi_{j_1, \dots, j_\ell}^{m_1, \dots, m_\ell} | v \rangle,
 \end{aligned}$$

where $\Psi_{i,j,m} = \sum_{1 \leq p \leq \ell_{i,j}-m} |v_{i,a_{i,j}+p}\rangle \langle u_{i,a_{i,j}+m+p}|$ and

$$\Psi_{j_1, \dots, j_\ell}^{m_1, \dots, m_\ell} = \Psi_{1,j_1,m_1} \Psi_{2,j_2,m_2} \cdots \Psi_{\ell,j_\ell,m_\ell}.$$

We may split the above summation, as before, into two parts corresponding to products containing only dominant eigenvalues to powers of free variables and those containing at least one subdominant eigenvalue to the power of a free variables. By Lemma 5, Jordan blocks corresponding to dominant eigenvalues have size 1×1 , that is, if $|\lambda_{t,j_t}| = 1$ for $t \in J_F$, then $\ell_{t,j_t} = 1$ and hence $m_t = 0$. So, we can write

$$\langle u | A_1^{k_1} A_2^{k_2} \cdots A_\ell^{k_\ell} | v \rangle = S_0 + S_1,$$

where

$$\begin{aligned}
 S_0 &= \sum_{\substack{j_1, \dots, j_\ell \mid j_q \in [1, n_q] \\ m_1, \dots, m_\ell \mid m_q \in [0, \ell_{q,j_q}-1] \\ \forall t \in J_F : |\lambda_{t,j_t}| = 1}} \left(\prod_{t \in J_F} \lambda_{t,j_t}^{k_t} \right) \Theta_{j_1, \dots, j_\ell}^{m_1, \dots, m_\ell}, \\
 S_1 &= \sum_{\substack{j_1, \dots, j_\ell \mid j_q \in [1, n_q] \\ m_1, \dots, m_\ell \mid m_q \in [0, \ell_{q,j_q}-1] \\ \exists t \in J_F : |\lambda_{t,j_t}| < 1}} \left(\prod_{t \in J_F} \lambda_{t,j_t}^{k_t-m_t} \binom{k_t}{m_t} \right) \Theta_{j_1, \dots, j_\ell}^{m_1, \dots, m_\ell} \quad \text{and} \\
 \Theta_{j_1, \dots, j_\ell}^{m_1, \dots, m_\ell} &= \left(\prod_{t \in J \setminus J_F} \lambda_{t,j_t}^{k_t-m_t} \binom{k_t}{m_t} \right) \langle u | \Psi_{j_1, \dots, j_\ell}^{m_1, \dots, m_\ell} | v \rangle.
 \end{aligned} \tag{5}$$

Note that k_t 's in the formula for $\Theta_{j_1, \dots, j_\ell}^{m_1, \dots, m_\ell}$ are fixed since $t \notin J_F$ and so A_t is not a free matrix. In other words, $\Theta_{j_1, \dots, j_\ell}^{m_1, \dots, m_\ell}$ does not depend on free variables k_i for $t \in J_F$. This also implies that S_0 assumes only finitely many different values as k_t with $t \in J_F$ vary since by Lemma 5 the dominant eigenvalues are roots of unity.

Again using the fact that Jordan blocks corresponding to the dominant eigenvalues have size 1×1 , we can rewrite the product inside the formula for S_1 from Eqn (5) as follows

$$\prod_{t \in J_F} \lambda_{t, j_t}^{k_t - m_t} \binom{k_t}{m_t} = \prod_{\substack{t \in J_F \\ |\lambda_{t, j_t}| = 1}} \lambda_{t, j_t}^{k_t} \cdot \prod_{\substack{t \in J_F \\ |\lambda_{t, j_t}| < 1}} \lambda_{t, j_t}^{k_t - m_t} \binom{k_t}{m_t}.$$

Suppose S_1 is not an empty sum since otherwise $S_1 = 0$. Then there exists $t \in J_F$ and $j_t \in [1, n_t]$ such that $|\lambda_{t, j_t}| < 1$. Let ρ be the maximum among such values, that is,

$$\rho = \max\{|\lambda_{t, j}| : t \in J_F, j \in [1, n_t] \text{ and } |\lambda_{t, j}| < 1\}.$$

Notice that every summand in S_1 has at least one $|\lambda_{t, j_t}| \leq \rho < 1$ with $t \in J_F$, and $|\lambda_{i, j}| \leq 1$ for all other $\lambda_{i, j}$. Also, $\binom{k_t}{m_t} \leq k_t^{m_t} \leq k_t^n$ since $m_t \leq n$. So every summand in S_1 can be estimated by the expression

$$C_1 \cdot \prod_{\substack{t \in J_F \\ |\lambda_{t, j_t}| < 1}} \rho^{k_t} k_t^n, \quad \text{where } C_1 \text{ is a computable constant.}$$

We have $\rho^k k^n \rightarrow 0$ when $k \rightarrow \infty$, and for any rational $\delta > 0$ we can compute C such that $\rho^k k^n < \delta$ for $k \geq C$. If in addition we assume that $0 < \delta < 1$ and that $k_t \geq C$ for all $t \in J_F$, then $|S_1| \leq C_1 n^{2\ell} \delta$.

Now, S_0 gives a finite number of limit values for $\langle u | A_1^{k_1} A_2^{k_2} \cdots A_\ell^{k_\ell} | v \rangle$. If λ is not equal to any of them, then choose a rational $0 < \delta < 1$ such that $\epsilon = C_1 n^{2\ell} \delta$ is less than half the minimal distance between λ and those limit values. Using this δ , we compute C as before. By definition of C , if all $k_t \geq C$ for $t \in J_F$, then the distance between $\langle u | A_1^{k_1} A_2^{k_2} \cdots A_\ell^{k_\ell} | v \rangle$ and one of the limit values of S_0 is less than ϵ . Thus $\langle u | A_1^{k_1} A_2^{k_2} \cdots A_\ell^{k_\ell} | v \rangle$ cannot be equal to λ when all $k_t \geq C$ for $t \in J_F$. Hence if $\lambda = \langle u | A_1^{k_1} A_2^{k_2} \cdots A_\ell^{k_\ell} | v \rangle$, then there is $t \in J_F$ such that $k_t < C$.

5 Other decidability results and NP-hardness

In this section we utilise Theorem 1 to obtain some related decidability results. The first of these combines Theorem 1 with a seminal result of Rabin and allows us to use our decidability result for cutpoint isolation to solve the emptiness problem for PFA on letter-bounded context-free languages when the cutpoint is isolated. We again highlight here that the emptiness problem is undecidable in general on letter-bounded languages, even when all matrices commute and the PFA is polynomially ambiguous [4].

► **Corollary 9.** *The emptiness problem is decidable for probabilistic finite automata on letter-bounded context-free languages when the cutpoint is isolated.*

Proof. A seminal result of Rabin [23] showed that given a n -state PFA \mathcal{P} acting on an alphabet Σ and *isolated* cutpoint $\lambda \in [0, 1]$ such that λ is isolated by $\epsilon > 0$ (i.e. $|\mathcal{P}(w) - \lambda| > \epsilon$ for all $w \in \Sigma^*$), then there exists a DFA \mathcal{D} such that $L_{< \lambda}(\mathcal{P}) = L(\mathcal{D})$, where $L(\mathcal{D})$ denotes the language accepted by the DFA \mathcal{D} . Moreover, Rabin showed that the number of states of \mathcal{D} is no more than $\left(1 + \frac{|F|}{\epsilon}\right)^{n-1}$ where F is the set of final states of \mathcal{P} .

22:14 Decidability of Cutpoint Isolation for PFA on Letter-Bounded Inputs

We note that the proof of Theorem 1 not only determines if a cutpoint is isolated but also determines an “isolation bound” $\epsilon > 0$ if it is isolated. In this case we can use Rabin’s result to construct an equivalent DFA $\mathcal{D}_<$ recognising $L_{<\lambda}(\mathcal{P})$. By inverting final and non final states of $\mathcal{D}_<$, we can construct \mathcal{D}_\geq which recognises $L_{\geq\lambda}(\mathcal{P})$. Finally we note that if λ is isolated then $L_{<\lambda}(\mathcal{P}) = L_{\leq\lambda}(\mathcal{P})$ and thus $\mathcal{D}_<$ and \mathcal{D}_\geq recognise the same languages as $L_{\leq\lambda}(\mathcal{P})$ and $L_{>\lambda}(\mathcal{P})$, respectively. Hence the emptiness problem is decidable. ◀

We now show that the value-1 problem for PFA on letter-bounded CFL inputs is decidable. This problem is undecidable for standard PFA but decidable for #-cyclic automata [13].

► **Corollary 10.** *The value-1 is decidable for probabilistic finite automata on letter-bounded context-free languages.*

Proof. This is trivial since the value-1 problem is equivalent to the isolation of the cutpoint 1 for a PFA [13]. ◀

Finally we note that Theorem 1 trivially allows us to determine if a given cutpoint is isolated for a PFA which is allowed a fixed maximum number of alternations between input letters (in any order), where each letter may be taken to an arbitrarily high power.

► **Corollary 11.** *Given a probabilistic finite automaton \mathcal{P} on alphabet $\Sigma = \{a_1, \dots, a_\ell\}$, cutpoint $\lambda \in [0, 1]$ and maximum number $k > 0$ of alternations between input letters, then determining if the cutpoint is isolated is decidable.*

Proof. We may apply Algorithm 1 on \mathcal{P} and λ with each language from the following (finite) set of letter-bounded languages $\Lambda = \{w_1^* w_2^* \dots w_k^* \mid w_i \in \Sigma\}$.

This defines the set of inputs where we alternate between the input letters a maximum of k times (analogous to how counter automata models are often studied with a maximum number of alternations between increasing and decreasing the counters). If any $L \in \Lambda$ on Algorithm 1 returns that the cutpoint is not isolated then λ is not isolated for \mathcal{P} with a maximum number of alternations k , otherwise the cutpoint is isolated. ◀

In the remainder of this section, we give a lower bound on the complexity of the cutpoint isolation problem for 3-state PFA on letter-bounded inputs (noting that the encoded PFA are polynomially rather than exponentially ambiguous).

► **Theorem 12.** *Cutpoint isolation is NP-hard for 3-state PFA on letter-bounded inputs.*

Proof. We use a reduction from the subset sum problem, defined thus: given a set of positive integers $S = \{x_1, x_2, \dots, x_k\} \subseteq \mathbb{N}$ and a natural number $T \in \mathbb{N}$, does there exist a subset $S' \subseteq S$ such that $\sum_{\ell \in S'} \ell = T$? This problem is well known to be NP-complete [12]. We define the set of matrices $M = \{A_i, B_i \mid 1 \leq i \leq k\} \subseteq \mathbb{Q}^{3 \times 3}$ in the following way:

$$A_i = \frac{1}{x_i + 1} \begin{pmatrix} 1 & x_i & 0 \\ 0 & 1 & x_i \\ 0 & 0 & x_i + 1 \end{pmatrix}, \quad B_i = \frac{1}{x_i + 1} \begin{pmatrix} 1 & 0 & x_i \\ 0 & 1 & x_i \\ 0 & 0 & x_i + 1 \end{pmatrix}$$

Note that A_i and B_i are thus row stochastic. Let $u = (1, 0, 0)^\top$ be the initial probability distribution, $v = (0, 1, 0)^\top$ be the final state vector and let $\mathcal{P} = (\langle u \mid, \{A_i, B_i\}, \mid v \rangle)$ be our PFA. We define the cutpoint $\lambda = \frac{T}{y}$, where $y = \sum_{j=1}^k (x_j + 1)$. Define letter-bounded language $\mathcal{L} = (a_1 \mid b_1)(a_2 \mid b_2) \dots (a_k \mid b_k) \subseteq a_1^* b_1^* a_2^* b_2^* \dots a_k^* b_k^*$ (thus \mathcal{L} is letter monotonic) and define a morphism $\varphi : \{a_i, b_i \mid 1 \leq i \leq k\}^* \rightarrow \{A_i, B_i \mid 1 \leq i \leq k\}^*$ in the natural way (e.g. the morphism induced by $\varphi(a_i) = A_i$ and $\varphi(b_i) = B_i$). Now, for a word $w = w_1 w_2 \dots w_k \in \mathcal{L}$,

note that $w_j \in \{a_j, b_j\}$ for $1 \leq j \leq k$. Define that $\mathbf{v}(a_i) = x_i$ and $\mathbf{v}(b_i) = 0$ and inductively extend to $\mathbf{v} : \Sigma^* \rightarrow \mathbb{N}$ by defining $\mathbf{v}(w_1 w_2 \cdots w_k) = \mathbf{v}(w_1) + \mathbf{v}(w_2 \cdots w_k)$ with $\mathbf{v}(\varepsilon) = 0$. In this case, we see that (due to the structure of A_i and B_i):

$$\langle u | \varphi(w_1 w_2 \cdots w_k) | v \rangle = \frac{\mathbf{v}(w)}{\prod_{j=1}^k (x_j + 1)}$$

Note of course that the factor $\frac{1}{\prod_{j=1}^k (x_j + 1)}$ is the same for any $w \in \mathcal{L}$.

Assume that there exists a solution to the subset sum problem, i.e., there exists $S' \subseteq S$ such that $\sum_{\ell \in S'} \ell = T$. Then consider word $w = w_1 w_2 \cdots w_k$ such that $w_j = a_j$ if $x_j \in S'$ and $w_j = b_j$ otherwise. In this case, $\sum_{i \in S'} x_i = \mathbf{v}(w)$ and thus $\langle u | \varphi(w_1 w_2 \cdots w_k) | v \rangle = \frac{T}{y} = \lambda$. If no solution exists, then for any word $w = w_1 w_2 \cdots w_k$, $|\mathbf{v}(w) - T| \geq 1$, and so $|\langle u | \varphi(w_1 w_2 \cdots w_k) | v \rangle - \lambda| > \frac{1}{\prod_{j=1}^k (x_j + 1)}$ and thus λ cannot be arbitrarily approximated.

Clearly the representation size of the PFA \mathcal{P} and λ are polynomial in the representation size of the subset sum problem instance and therefore we are done. ◀

6 Conclusion

In this work we showed that the cutpoint isolation problem is decidable for PFA when the input words are constrained to come from a letter-bounded context-free language, even for exponentially ambiguous PFA. This is in contrast to the situation for the (strict) emptiness problem and the injectivity problem, which are *undecidable* even over more restricted PFA for which all matrices commute, the PFA is polynomially ambiguous and the input words are over a simple letter-bounded language $a_1^* \cdots a_\ell^*$. We show that if the cutpoint is isolated for words over the input language, then the emptiness problem becomes decidable. We also show that the value-1 problem is decidable for these restricted input words.

It would be interesting to determine the complexity of the cutpoint isolation problem more precisely. We show an NP-hard lower bound in Theorem 12. The algorithm we provide may belong to NP, however there are some issues with showing this upper bound, namely that the number of limits of a stochastic matrix may be exponential in its dimension and the value of constant C from Proposition 7 may be exponential in terms of the bit size of the matrices, making the verification stage of an NP algorithm difficult to achieve. Extending the results to more general bounded languages would also be an interesting future work.

References

- 1 Saugata Basu, Richard Pollack, and Marie-Françoise Roy. *Algorithms in real algebraic geometry*, volume 10 of *Algorithms and Computation in Mathematics*. Springer-Verlag, Berlin, second edition, 2006.
- 2 P. C. Bell, S. Chen, and L. M. Jackson. Freeness properties of weighted and probabilistic automata over bounded languages. *Information and Computation*, 269, 2019.
- 3 P. C. Bell, V. Halava, and M. Hirvensalo. Decision problems for probabilistic finite automata on bounded languages. *Fundamenta Informaticae*, 123(1):1–14, 2012.
- 4 Paul C. Bell. Polynomially ambiguous probabilistic automata on restricted languages. In *International Colloquium on Automata, Languages, and Programming (ICALP'19)*, volume 132 of *LIPICs*, pages 105:1–105:14, 2019. doi:10.4230/LIPICs.ICALP.2019.105.
- 5 A. Bertoni, G. Mauri, and M. Torelli. Some recursively unsolvable problems relating to isolated cutpoints in probabilistic automata. In *Automata, Languages and Programming*, volume 52, pages 87–94, 1977.

- 6 V. Blondel and V. Canterini. Undecidable problems for probabilistic automata of fixed dimension. *Theory of Computing Systems*, 36:231–245, 2003.
- 7 Jin-yi Cai. Computing Jordan normal forms exactly for commuting matrices in polynomial time. *Int. J. Found. Comput. Sci.*, 5(3/4):293–302, 1994. doi:10.1142/S0129054194000165.
- 8 Henri Cohen. *A course in computational algebraic number theory*, volume 138 of *Graduate Texts in Mathematics*. Springer-Verlag, Berlin, 1993. doi:10.1007/978-3-662-02945-9.
- 9 L. Daviaud, M. Jurdzinski, R. Lazic, F. Mazowiecki, G. A. Pérez, and J. Worrell. When is containment decidable for probabilistic automata? In *International Colloquium on Automata, Languages, and Programming (ICALP'18)*, volume 107 of *LIPICs*, pages 121:1–121:14, 2018. doi:10.4230/LIPICs.ICALP.2018.121.
- 10 N. Fijalkow, C. Riveros, and J. Worrell. Probabilistic automata of bounded ambiguity. In *28th International Conference on Concurrency Theory (CONCUR)*, pages 19:1–19:14, 2017.
- 11 S. Friedland. *Matrices: Algebra, Analysis and Applications*. World Scientific Publishing Company Pte Limited, 2015. URL: <https://books.google.co.uk/books?id=y8fACwAAQBAJ>.
- 12 M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Co. New York, NY, USA, 1979.
- 13 H. Gimbert and Y. Oualhadj. Probabilistic automata on finite words: decidable and undecidable problems. In *International Colloquium on Automata, Languages and Programming (ICALP'10)*, volume 2, pages 527–538, 2010.
- 14 S. Ginsburg. *The Mathematical Theory of Context Free Languages*. McGraw-Hill, 1966.
- 15 Seymour Ginsburg and Edwin Spanier. Semigroups, Presburger formulas, and languages. *Pacific journal of Mathematics*, 16(2):285–296, 1966.
- 16 O. Goldreich. On promise problems: a survey. In *Essays in Memory of Shimon Even*, volume 3895 of *Lecture Notes in Computer Science*, pages 254–290. Springer-Verlag, 2006.
- 17 V. Halava, T. Harju, M. Hirvensalo, and J. Karhumäki. Skolem’s problem — on the border between decidability and undecidability. In *TUCS Technical Report Number 683*, 2005.
- 18 M. Hirvensalo. Improved undecidability results on the emptiness problem of probabilistic and quantum cut-point languages. *SOFSEM 2007: Theory and Practice of Computer Science, Lecture Notes in Computer Science*, 4362:309–319, 2007.
- 19 M. Mignotte. Some useful bounds. In *Computer algebra*, pages 259–263. Springer, Vienna, 1983.
- 20 V. Y. Pan. Optimal and nearly optimal algorithms for approximating polynomial zeros. *Comput. Math. Appl.*, 31(12):97–138, 1996. doi:10.1016/0898-1221(96)00080-6.
- 21 R. J. Parikh. On context-free languages. *Journal of the ACM (JACM)*, 13(4):570–581, 1966.
- 22 A. Paz. *Introduction to Probabilistic Automata*. Academic Press, 1971.
- 23 M. O. Rabin. Probabilistic automata. *Information and Control*, 6:230–245, 1963.
- 24 J. C. Rosales and P. A. García-Sánchez. *Numerical semigroups*, volume 20 of *Developments in Mathematics*. Springer, New York, 2009. doi:10.1007/978-1-4419-0160-6.
- 25 P. Turakainen. Generalized automata and stochastic languages. *Proceedings of the American Mathematical Society*, 21:303–309, 1969.
- 26 A. Weber and H. Seidl. On the degree of ambiguity of finite automata. *Theoretical Computer Science*, 88(2):325–349, 1991.