Laminations as finite state automata

Eduardo Ventilari Sodré

May 5, 2026

1 ω-automata and shift invariance

A *(deterministic fail) ω-automaton* is a tuple $A = (Q, \Sigma, \delta, s, f)$ consisting of:

- A finite set of *states* Q;
- a finite set of symbols Σ, the *alphabet* of the language;
- A *transition function* $\delta : Q \times \Sigma \to Q$, which determines how one state transitions to others based on the input $a \in \Sigma$;
- A specified *start state* $s \in Q$;
- A specified *fail state* $f \in Q$, such that for all $a \in \Sigma$, $\delta(f, a) = f$.

Consider the set $\Sigma^{<\omega}$ of all finite words with letters in Σ, and similarly Σ^{ω} the set of (one-sided) infinite words in Σ. The finite words form an associative, non-commutative monoid under concatenation, and we may similarly concate a finite word on the left with an infinite word on the right. An *ω-language* L is any subset of Σ^{ω}.

We say that an ω-automaton A *recognizes* the ω-language L if L consists of exactly the infinte words for which A does not ever enter its fail state starting from its start state. More precisely, we may extend the definition of δ to a transition function $\overline{\delta} : Q \times \Sigma^{<\omega} \to Q$ that "reads off" each letter of a word to decide the transition from the current state to the next. Recursively, if a finite word is of the form aw for $a \in \Sigma$, then

$$\overline{\delta}(q, aw) = \overline{\delta}(\delta(q, a), w),$$

where by convention if ε is the empty word, $\overline{\delta}(q, \varepsilon) = q$. (Often we will denote $\overline{\delta}$ also as δ.) Therefore A recognizes L if L consists exactly of those infinite words $w = a_1a_2\dots$ such that, for all truncations $w[1,n] = a_1a_2\dots a_n$ of w, we have $f \neq \overline{\delta}(s, w[1,n])$. Finally, we say that $L \subseteq \Sigma^\omega$ is an *ω-regular language* if it is recognized by some ω-automaton. Intuitively, the automaton either stops at at some point to decide if a given word fails to be in the language, or continues on forever, since it cannot decide if the word is in L without reading all of its letters to check if it never enters its fail state.

There are several more general definitions for ω-automata, with distinct acceptance conditions other than never entering a given fail state. However, this definition is sufficient for our purposes, and a fail ω-automaton will still be an ω-automaton for any of the other common definitions in the literature. We may also extend the definition to allow for some finite collection of distinct fail states $f_1, \dots, f_k$, each of which transitions only to itself, and we do so without harm when necessary.

Let $\sigma : \Sigma^\omega \to \Sigma^\omega$ be the *shift map*, that is, if $w \in \Sigma^\omega$ is of the form $w = a_1a_2a_3\dots$ for $a_i \in \Sigma$, then

$$\sigma(a_1a_2a_3\dots) = a_2a_3a_4\dots$$

As Σ is finite, the set Σ^ω has a natural Tychonoff topology, with respect to which it is homeomorphic to a Cantor set and σ is continuous and surjective. We say that an ω-language is *shift invariant* if $\sigma(L) \subseteq L$. We are interested in those ω-regular languages that are also shift invariant.

Proposition 1.1. *Let L be an ω-regular language on a finite alphabet Σ, and let $\sigma : \Sigma^\omega \to \Sigma^\omega$ be the shift on one-sided infinite words in Σ. Let*

$$L_\sigma := \bigcap_{i \geq 0} \sigma^{-i}(L) = \{w \in \Sigma^\omega \mid \sigma^i(w) \in L \ \forall i \geq 0\}.$$

Then L_σ is a shift invariant ω-regular language.

Proof. Let A be an ω-automaton that recognizes L, with set of states Q, transition function $\delta : Q \times \Sigma \to Q$, start state s and fail state f. By modifying A (adding at most one state) we may assume that no state transitions to s. We construct an ω-automaton A_σ that recognizes L_σ intituively as the same as A, but with a non-deterministic transition from s to s. This corresponds to "ignoring" the first input, and by ignoring any finite amount of initial inputs the automaton verifies if the rest of the word is also in the language. As this happens for all possible shifts, the automaton recognizes L_σ.

To describe A_σ as deterministic ω-automaton, we first consider its set of states to be $2^Q \setminus \{\emptyset\}$. The transitions are as follows:

1. if $S \in 2^Q$ is such that $f \in S$, then all inputs transition S to $\{s, f\}$.

2. If $f, s \notin S = \{q_1, \ldots, q_m\}$, then S transitions to
$$\{\delta(q_1, a), \ldots, \delta(q_m, a)\}$$
with the input a.

3. If $f \notin S$ and $s \in S = \{q_1, \ldots, q_m\}$, then S transitions to
$$\{\delta(q_1, a), \ldots, \delta(q_m, a), s\}$$
with the input a.

We stipulate that $\{s\}$ is the start state for A_σ. Hence only rules 1 and 3 are relevant, as any reachable state will contain s. If w is a finite word of length n, then $\delta(\{s\}, w)$ is exactly

$$\{\delta(s, w), \delta(s, \sigma(w)), \ldots, \delta(s, \sigma^{n-1}(w)), \delta(s, \sigma^n(w))\}$$

where we extend σ to finite words by erasing the first letter, and by convention $\delta(s, \sigma^n(w)) = \delta(s, \epsilon) = s$. This implies that $\delta(\{s\}, w)$ contains f if and only if the fail state is reached for some shift of w. From this, we readily conclude that A_σ recognizes L_σ. That L_σ is shift invariant is immediate from the definition. $\square$

2 Laminations

Thurston's landmark paper [**?**] describes the theory of (gemeotric) invariant laminations, which serve as a combinatorial model for polynomial dynamics; we restrict ourselves to the quadratic case. There are two main ways to create invariant laminations, and they are "minimal" in a specific sense in most cases.

For the first, let $\theta \in [0, 1)$ be a real number, and l be the diameter of the closed unit disk $\overline{\mathbb{D}}$ connecting the boundary points at angles $\theta/2$ and $(1+\theta)/2$, measured in turns of the circle (henceforth we will indeitfy the angle with the point on $S^1 = \partial\mathbb{D}$.) This trivially describes a forward invariant quadratic lamination, with the leaf collapsing to the "degenerate leaf" at the point $\theta \in S^1$. By Proposition $\cdot$ in [**?**], there is an invariant lamination $\mathcal{L}_\theta$ that contains this forward invariant one. The construction of $\mathcal{L}_\theta$ is recursive, made by adjoining at least 2 "preimages" for each leaf such that, out of all possibilities, they do not cross over leaves in previous steps of the construction. This has some ambiguity only in the case where θ is periodic, where a leaf may have θ as one if its endpoints, and so we add all four possible preimages for it, constructing two symmetric collapsing triangles. By taking the union of leaves at each step of the construction and taking the closure of all of them we obtain $\mathcal{L}_\theta$.

There is a simple combinatorial description of which leaves belong to the lamination $\mathcal{L}_\theta$. Given $\alpha \in S^1$, we describe its *kneading sequence* $\nu_\theta(\alpha) = \nu_0\nu_1\nu_2\dots$ as follows. Let $\sigma_2 : S^1 \to S^1$ be the doubling map $\alpha \mapsto 2\alpha$, and we use the shorthand notation $\alpha_i \coloneqq \sigma_2^i(\alpha)$. We define $\nu_i = 1$ if $\alpha_i \in (\theta/2, (1+\theta)/2)$, the open arc from $\theta/2$ to $(1+\theta)/2$ following counterclockwise from 0, $\nu_i = 0$ if $\alpha_i \in ((1+\theta)/2, \theta/2)$, and $\nu_i = \star$ if $\alpha_i = \theta/2$ or $(1+\theta)/2$. Note that except for when $\theta = 0$, we always have $\nu_0 = 1$.

Proposition 2.1. *The leaf l between α and β is in $\mathcal{L}_\theta$ if and only if $\nu_\theta(\alpha) = \nu_\theta(\beta)$.*

Our main goal is to describe these geometric quadratic laminations in terms of shift invariant ω-automata.

3 Comparison automata

Let $\theta \in [0, 1)$ be a fixed positive rational number. With respect to its binary expansion

$$\theta = 0.\theta_1\theta_2\theta_3\dots = \sum_{n=1}^{\infty} \frac{\theta_i}{2^i},$$

there are two possibilities; either θ is a dyadic rational of the form $p/2^n$ for integers $n \geq 0$ and $0 \leq p < 2^n$, or it isn't. In the former case, there are exactly two binary expansions for θ, and we choose the one that terminates, so that

$$\theta = 0.\theta_1\theta_2\dots\theta_k.$$

In the latter, its binary expansion is eventually periodic, and

$$\theta = 0.\theta_1\dots\theta_l\overline{\theta_{l+1}\dots\theta_{l+k}}.$$

It is possible to describe an ω-automaton A which takes as an input a real number $\alpha \in [0, 1)$, in the form of its binary expansion, and decides if $\alpha = \theta$, $\alpha > \theta$ or $\alpha < \theta$. We assume that the (terminating or periodic) binary expansion of θ is known, but α may be arbitrary, where the automaton reads the digits of α one by one.

We may construct A so that the number of its states is linear with respect to the relevant digits of θ, where intuitively the state keeps track of both the current state of the binary expansion of θ (which digit we are in the preperiod of period, or of the terminating expansion), and what is the current state of comparison (if the binary expansions of α and θ are equal so far, or the automata has witnessed $\alpha > \theta$ or $\alpha < \theta$). Note that it is possible that α as an input for a binary expansion has terminating period $0.\alpha_1 \dots \alpha_k \overline{1}$, so that just witnessing a position in which the digits of α and θ is not sufficient to witness the strict inequality. This, however, can still be encoded within the states of the automaton. Here we also use the extended definition of a deterministic fail ω-automaton, allowing for multiple fail states, which in this case witness whether $\alpha > \theta$ or $\alpha < \theta$.

4 ω-automata for leaves

From the comparison automaton constructed before, for a fixed $\theta \in \mathbb{Q} \cap [0, 1)$ one may imagine the possibility of constructing an ω-automaton that takes as an input a real number $\alpha \in [0, 1)$ in the form of its binary expansion, and outputs:

$$\begin{cases} 1, & \text{if } \alpha \in (\theta/2, (1+\theta)/2); \\ 0, & \text{if } \alpha \in (1+\theta)/2, \theta/2); \\ \star, & \text{if } \alpha \in \{\theta/2, (1+\theta)/2\}, \end{cases}$$

that is, the first digit of its kneading sequence $\nu_\theta(\alpha)$. Strictly speaking, this is not necessarily done in finite time, as equalities would require the automaton to go on forever without witnessing a strict inequality. However, given Proposition 1.1 and Proposition 2.1, if we extend our alphabet to $\Sigma = \{0, 1\}^2$ and adjust our automata to simultaneously read off the binary expansion of two real numbers $\alpha, \beta \in [0, 1)$, we obtain:

Proposition 4.1. *Given $\theta \in \mathbb{Q} \cap [0, 1)$, there exists a deterministic ω-automaton that recoginzes exactly when a leaf l, determined by its endpoints $l \cap S^1 = \{\alpha, \beta\}$, is in the lamination $\mathcal{L}_\theta$ or not.*

As a remark, even if α or β are dyadic rationals, the automaton can still recognize the leaf no matter which binary expansion is chosen.

Let L_θ be the ω-regular language recognized by this automaton. From its construction, L_θ is shift invariant. A natural question to ask, in the context of symbolic dynamics, is if L_θ is a subshift of finite type on Σ^ω, or if not, how much does it fail to be one. To this end, we introduce a generalization of subshifts of finite type. Given a set of symbols Σ and $n \geq 1$, an *n-subshift of finite type* is a subset $A \subseteq \Sigma^\omega$ of infinite words such that there exists a set $S \subseteq \Sigma^n$ of *admissible blocks* of length n such that

$$w \in A \iff w[i, i+n-1] \in S, \ \forall i \geq 1.$$

This says that A consists of all infinite words such that any length n segment is admissible, that is, corresponds to one of the allowed blocks, represented by the set S. Note that $n = 2$ corresponds to a classical subshift of finite type, and $n = 1$ is merely those sequences consisting only of a specific subset of symbols. We may denote $A = \Sigma(S)$.

As an initial example, let $\theta = 1/2$. The lamination $\mathcal{L}_{1/2}$ is the vertical lamination, where a leaf $\overline{\alpha\beta}$ is in $\mathcal{L}_\theta$ if and only if $a + \beta = 1$. This can be expressed as a 1-subshift of finite type, since the condition $\alpha + \beta = 1$, translated to binary sequences, just means that the binary digits for α and β at the same position must always be distinct, either $(0, 1)$ or $(1, 0)$.

This feels wrong: choice of binary expansion gives different symbols, possibly equal...

Now consider any rational $\theta \in (1/3, 2/3)$, so that $\theta/2 \in (1/6, 1/3)$ and therefore the diamater does not separate the angles $1/3$ and $2/3$. As they form an orbit of period 2 under doubling, they naturally have the same kneading sequence and the leaf connecting them is in $\mathcal{L}_\theta$ (as we pass to the closure of the inductive construction). In particular, the word

$$\begin{bmatrix} 0.\overline{01} \\ 0.\overline{10} \end{bmatrix}$$

is in L_θ. As the laminations $\mathcal{L}_\theta$ are also backward invariant, we may take preimages of leaves also in the lamination, which, for a given pair of binary expansions, corresponds to inserting a zero or one at the start of one, and a zero or one at the start of the other. For example, the word for the leaf between $1/6$ and $5/6$ is

$$\begin{bmatrix} 0.0\overline{01} \\ 0.1\overline{10} \end{bmatrix},$$

and more generally L_θ will contain infinitely many words such that the binary expansions are eventually periodic, of the form

$$\begin{bmatrix} 0.u\overline{01} \\ 0.v\overline{10} \end{bmatrix}$$

where u and v are finite words of the same length.

For $k \geq 1$, Let $S_i^n \subseteq \Sigma^n$ be the set of words of length n that occur as $w[i, i+n-1]$ for words w in L_θ. From the forward invariance of leaves, which implies shift invariance of L_θ, we have

$$S_1^n \supseteq S_2^n \supseteq S_3^n \supseteq \cdots,$$

and from backward invariance, we also get

$$S_1^n \subseteq S_2^n \subseteq S_3^n \subseteq \cdots.$$

Hence all sets S_i^n are equal for fixed n, and we denote them by S^n.

If L_θ were a n-subshift of finite type $\Sigma(S)$ for some $S \subseteq \Sigma^n$, then necessarily $S^n \subseteq S$, as there may in principle be words in S which do not appear in any actually realizable infinite word in $\Sigma(S)$. However, we may remove these extraneuous unrealizable words and obtain the same subset of Σ^ω, so without loss of generality we may assume that $S = S^n$. Also without loss of generality, we may assume that n is minimal with respect to the property that L_θ is an n-subshift of finite type, since any n-subshift of finite type is also naturally an $(n+1)$-subshift of finite type; and that S is a minimal set of words which generates the same subshift, that is, minimal such that $\Sigma(S) = L_\theta$.

Proposition 4.2. *For $\theta \in (1/3, 2/3)$ and $\theta \neq 1/2$, L_θ is not a 2-subshift of finite type.*

Proof. From the previous analysis, we see that

$$\begin{bmatrix} 0.\overline{01} \\ 0.\overline{10} \end{bmatrix}, \quad \begin{bmatrix} 0.0\overline{01} \\ 0.1\overline{10} \end{bmatrix}$$

are words in L_θ, and if $L_\theta = \Sigma(S)$ for $S \subseteq \Sigma^2$, then necessarily

$$\begin{bmatrix} 01 \\ 10 \end{bmatrix}, \begin{bmatrix} 10 \\ 01 \end{bmatrix}, \begin{bmatrix} 00 \\ 11 \end{bmatrix} \in S,$$

as these are all length 2 words that occur as length 2-blocks for the above infinite words in L_θ. But then the word

$$\begin{bmatrix} 01000\ldots \\ 10111\ldots \end{bmatrix} = \begin{bmatrix} 01\overline{0} \\ 10\overline{1} \end{bmatrix}$$

would also be in L_θ, corresponding to the diameter between 1/4 and 3/4. This contradicts the assumptions made. ☐

Proposition 4.3. *L_θ is not a 3-subshift of finite type.*

Proof. The presence of the leaves between 5/12 and 7/12 also would produce the diameter 1/4, 3/4. ☐

Proposition 4.4. *If $\theta \in (\frac{1}{2} - \frac{1}{3\cdot 2^n}, \frac{1}{2} + \frac{1}{3\cdot 2^n})$ and $\theta \neq 1/2$, then L_θ is not an n-subshift of finite type.*

Proof. By taking preimages of the leaf $\{1/3, 2/3\}$, we may show that the leaf $\{\frac{1}{4} - \frac{1}{3\cdot 2^{n+1}}, \frac{3}{4} + \frac{1}{3\cdot 2^{n+1}}\}$ will be in $\mathcal{L}_\theta$. It's corresponding word is

$$\begin{bmatrix} 0.\overbrace{0100\ldots0}^{n}\overline{01} \\ 0.1011\ldots1\overline{10} \end{bmatrix},$$

and using the blocks from the word corresponding to this leaf and

$$\begin{bmatrix} 00\ldots0 \\ 11\ldots1 \end{bmatrix}$$

which is an admissible block for any L_θ, we can still show that the leaf $\{1/4, 3/4\}$ would be in the lamination. $\square$