Tree Shift Complexity

Karl Petersen and Ibrahim Salama

University of North Carolina at Chapel Hill
and North Carolina Central University

Carolina Dynamics Symposium
April 14, 2018

Abstract

A tree shift consists of all labelings of the vertices of a tree by elements of a finite alphabet that omit all of a prescribed (possibly infinite) set of finite patterns.

Abstract

A tree shift consists of all labelings of the vertices of a tree by elements of a finite alphabet that omit all of a prescribed (possibly infinite) set of finite patterns.

These are exactly the closed sets that are invariant under the shifts associated with the tree.

Abstract

A tree shift consists of all labelings of the vertices of a tree by elements of a finite alphabet that omit all of a prescribed (possibly infinite) set of finite patterns.

These are exactly the closed sets that are invariant under the shifts associated with the tree.

Introduced by Aubrun and Béal and studied also by Ban and Chang, they share properties of the familiar one-dimensional subshifts of symbolic dynamics

Abstract

A tree shift consists of all labelings of the vertices of a tree by elements of a finite alphabet that omit all of a prescribed (possibly infinite) set of finite patterns.

These are exactly the closed sets that are invariant under the shifts associated with the tree.

Introduced by Aubrun and Béal and studied also by Ban and Chang, they share properties of the familiar one-dimensional subshifts of symbolic dynamics
while preserving a directional aspect that may make them easier to analyze than higher-dimensional subshifts, where questions of undecidability and computability arise.

We study the complexity function of a tree shift, which counts as a function of n the number of different labelings of a shape of size n.

We study the complexity function of a tree shift, which counts as a function of n the number of different labelings of a shape of size n.

We define the entropy in a different way than Ban and Chang, prove that the limit in the definition exists, and show that the entropy of a tree shift determined by adjacency constraints dominates the entropy of the associated one-dimensional subshift.

Dyadic tree: Σ^*

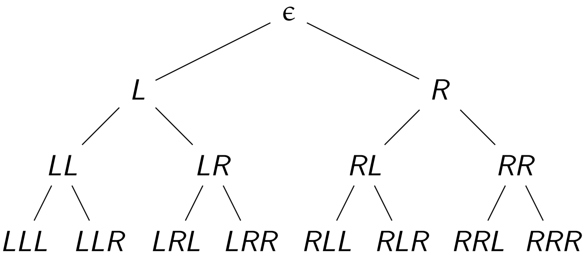

Labeled tree: $\tau : \Sigma^* \to A$

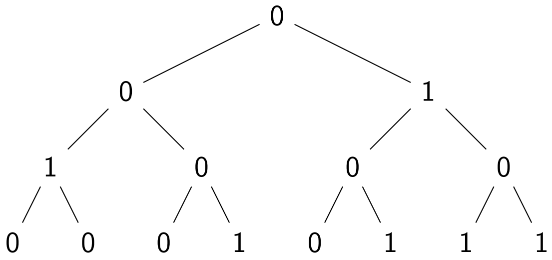

Topological entropy

The *complexity function* of a labeled tree τ is $p_\tau(n) =$ number of different labelings of subtrees Δ_n of height n seen in τ.

Topological entropy

The *complexity function* of a labeled tree τ is $p_\tau(n) =$ number of different labelings of subtrees Δ_n of height n seen in τ.

Theorem

For any labeled tree τ, the limit

$$h = h(\tau) = \lim_{n\to\infty} \frac{\log p_\tau(n)}{2^{n+1} - 1}$$

exists.

Topological entropy

The *complexity function* of a labeled tree τ is $p_\tau(n) =$ number of different labelings of subtrees Δ_n of height n seen in τ.

Theorem

For any labeled tree τ, the limit

$$h = h(\tau) = \lim_{n\to\infty} \frac{\log p_\tau(n)}{2^{n+1} - 1}$$

exists.

Proof: We do not have subadditivity, but the basic strategy still works.

Nearest neighbor k-tree shifts labeled by d symbols

Consider a dyadic tree shift with vertices labeled from a finite alphabet $A = \{a_1, \ldots, a_d\}$ with 1-step finite type restrictions given by a $0, 1$ matrix M indexed by the elements of A: adjacent nodes in the tree are allowed to have labels i for the first (closer to the root) and j for the second if and only if $M_{ij} = 1$.

Nearest neighbor k-tree shifts labeled by d symbols

Consider a dyadic tree shift with vertices labeled from a finite alphabet $A = \{a_1, \ldots, a_d\}$ with 1-step finite type restrictions given by a $0, 1$ matrix M indexed by the elements of A: adjacent nodes in the tree are allowed to have labels i for the first (closer to the root) and j for the second if and only if $M_{ij} = 1$.

We compare the entropy h of the tree shift with that of the corresponding one-dimensional shift of finite type.

Nearest neighbor k-tree shifts labeled by d symbols

Consider a dyadic tree shift with vertices labeled from a finite alphabet $A = \{a_1, \dots, a_d\}$ with 1-step finite type restrictions given by a $0, 1$ matrix M indexed by the elements of A: adjacent nodes in the tree are allowed to have labels i for the first (closer to the root) and j for the second if and only if $M_{ij} = 1$.

We compare the entropy h of the tree shift with that of the corresponding one-dimensional shift of finite type.

For the golden mean systems (no adjacent nodes or entries have the same label 1), $p_\tau(n)$ is the number of independent sets in a tree of height n,

Nearest neighbor k-tree shifts labeled by d symbols

Consider a dyadic tree shift with vertices labeled from a finite alphabet $A = \{a_1, \dots, a_d\}$ with 1-step finite type restrictions given by a $0, 1$ matrix M indexed by the elements of A: adjacent nodes in the tree are allowed to have labels i for the first (closer to the root) and j for the second if and only if $M_{ij} = 1$.

We compare the entropy h of the tree shift with that of the corresponding one-dimensional shift of finite type.

For the golden mean systems (no adjacent nodes or entries have the same label 1), $p_\tau(n)$ is the number of independent sets in a tree of height n,

a more complicated sequence than the Fibonacci numbers: $2, 5, 41, 2306 \dots$ (A076725).

Entropy comparison

Theorem

Let M be an irreducible d-dimensional $0, 1$ matrix, Σ_M the corresponding shift of finite type, and X_M the corresponding tree shift, labeled by elements of the alphabet A, with $|A| = d$, subject to the adjacency restrictions given by M.

Entropy comparison

Theorem

Let M be an irreducible d-dimensional $0, 1$ matrix, Σ_M the corresponding shift of finite type, and X_M the corresponding tree shift, labeled by elements of the alphabet A, with $|A| = d$, subject to the adjacency restrictions given by M.

Then the topological entropy of Σ_M is less than or equal to the topological entropy of X_M: $h_{\text{top}}(\Sigma_M) \leqslant h$.

Sketch of the proof

Let v denote the positive Perron-Frobenius left eigenvector of M normalized so that $\sum v_i = 1$, and let $\lambda > 0$ denote the maximum eigenvalue of M.

Sketch of the proof

Let v denote the positive Perron-Frobenius left eigenvector of M normalized so that $\sum v_i = 1$, and let $\lambda > 0$ denote the maximum eigenvalue of M.

For each $n = 0, 1, \ldots$ and $i = i, \ldots, |A|$, denote by $x(n) = (x_i(n)), i = 1, \ldots, |A|$, the vector that gives for each symbol $i \in A$ the number of trees of height n labeled according to the transitions allowed by M that have the symbol i at the root.

Sketch of the proof

Let v denote the positive Perron-Frobenius left eigenvector of M normalized so that $\sum v_i = 1$, and let $\lambda > 0$ denote the maximum eigenvalue of M.

For each $n = 0, 1, \dots$ and $i = i, \dots, |A|$, denote by $x(n) = (x_i(n)), i = 1, \dots, |A|$, the vector that gives for each symbol $i \in A$ the number of trees of height n labeled according to the transitions allowed by M that have the symbol i at the root.

Considering the symbols that can follow each symbol i in the last row of a labeling of Δ_n, and that they can be assigned in independent pairs to the nodes below, shows that these vectors satisfy the recurrence

Sketch of the proof

Let v denote the positive Perron-Frobenius left eigenvector of M normalized so that $\sum v_i = 1$, and let $\lambda > 0$ denote the maximum eigenvalue of M.

For each $n = 0, 1, \ldots$ and $i = i, \ldots, |A|$, denote by $x(n) = (x_i(n)), i = 1, \ldots, |A|$, the vector that gives for each symbol $i \in A$ the number of trees of height n labeled according to the transitions allowed by M that have the symbol i at the root.

Considering the symbols that can follow each symbol i in the last row of a labeling of Δ_n, and that they can be assigned in independent pairs to the nodes below, shows that these vectors satisfy the recurrence

$$x_i(0) = 1, \quad x_i(n+1) = (Mx(n))_i^2 \text{ for all } i = 1, \ldots, d, \text{ all } n \geqslant 0.$$

Proof continued

Denote by $\mathbf{1}$ the vector $(1, 1, \ldots, 1) \in \mathbb{R}^d$. We claim that

$$x(n) \cdot v \geqslant \lambda^{2^{n+1}-2} v \cdot \mathbf{1} \quad \text{for all } n \geqslant 0.$$

Proof continued

Denote by $\mathbf{1}$ the vector $(1, 1, \ldots, 1) \in \mathbb{R}^d$. We claim that

$$x(n) \cdot v \geqslant \lambda^{2^{n+1}-2} v \cdot \mathbf{1} \quad \text{for all } n \geqslant 0.$$

Since M is irreducible and all entries of v are positive, $x(n) \cdot v$ and $x(n) \cdot \mathbf{1}$ grow at the same superexponential rate, so the result would follow.

Proof continued

Denote by 1 the vector $(1, 1, \ldots, 1) \in \mathbb{R}^d$. We claim that

$$x(n) \cdot v \geqslant \lambda^{2^{n+1}-2} v \cdot 1 \quad \text{for all } n \geqslant 0.$$

Since M is irreducible and all entries of v are positive, $x(n) \cdot v$ and $x(n) \cdot 1$ grow at the same superexponential rate, so the result would follow.

We make an argument by induction. For $n = 0$ we have

$$x(0) \cdot v = \sum_i v_i = v \cdot 1.$$

Assuming that the inequality holds at stage n and using the inequality $\mathbb{E}(X^2) \geqslant [\mathbb{E}(X)]^2$ on the random variable $X_i = [Mx(n)]_i$ with discrete probabilities v_i, we have

Assuming that the inequality holds at stage n and using the inequality $\mathbb{E}(X^2) \geqslant [\mathbb{E}(X)]^2$ on the random variable $X_i = [Mx(n)]_i$ with discrete probabilities v_i, we have

$$\sum_i x_i(n+1)v_i = \sum_i (Mx(n))_i^2 v_i \geqslant \left[\sum_i Mx(n)_i v_i\right]^2$$

$$= \left[\sum_i x(n)_i (vM)_i\right]^2 = \left[\sum_i x(n)_i \lambda v_i\right]^2 = [\lambda x(n) \cdot v]^2$$

$$\geqslant \left[\lambda^{2^{n+1}-2} \lambda v \cdot 1\right]^2$$

$$= \lambda^{2^{n+2}-2} v \cdot 1.$$

Name	Matrix$=M$	$h_{\text{top}}(\Sigma_M)$	h (est)	U
Γ	$11, 10$	.481	.509	.721
X_0	$010, 101, 101$	.481	.509	.722
X_3	$011, 111, 101$	.81	.846	1.104
X_5	$110, 011, 101$	.693	.693	.693
X_{10}	$011, 111, 100$	.693	.774	1.242
X_{11}	$111, 100, 100$	.693	.763	1.04
A_1	$110, 101, 001$	.481	.611	∞
A_2	$110, 011, 010$	.481	.575	.962

Table: Estimates of some tree shift entropies

Asymptotics of ratios

For each $i = 0, 1, \ldots, d-1$ and $n \geqslant 0$ let $x_i(n)$ denote the number of blocks of height n that appear in Z_M and have the symbol i at the root. Let $x(n) = (x_0(n), \ldots, x_{d-1}(n)), |x(n)| = x_0(n) + \cdots + x_{d-1}(n)$, and $r(n) = r_{x(n)} = x(n)/|x(n)|$.

Asymptotics of ratios

For each $i = 0, 1, \ldots, d-1$ and $n \geqslant 0$ let $x_i(n)$ denote the number of blocks of height n that appear in Z_M and have the symbol i at the root. Let
$x(n) = (x_0(n), \ldots, x_{d-1}(n)), |x(n)| = x_0(n) + \cdots + x_{d-1}(n)$, and
$r(n) = r_{x(n)} = x(n)/|x(n)|$.

We have the recurrence

$$x_i(n+1) = (Mx(n))_i^2 \quad \text{for all } i = 0, \ldots, d-1, \text{ all } n \geqslant 0.$$

Asymptotics of ratios

For each $i = 0, 1, \dots, d-1$ and $n \geqslant 0$ let $x_i(n)$ denote the number of blocks of height n that appear in Z_M and have the symbol i at the root. Let
$x(n) = (x_0(n), \dots, x_{d-1}(n)), |x(n)| = x_0(n) + \cdots + x_{d-1}(n)$, and
$r(n) = r_{x(n)} = x(n)/|x(n)|$.

We have the recurrence

$$x_i(n+1) = (Mx(n))_i^2 \quad \text{for all } i = 0, \dots, d-1, \text{ all } n \geqslant 0.$$

Define $T : \mathbb{R}^d \setminus \{0\} \to \mathbb{R}^d \setminus \{0\}$ by

$$(Tx)_i = \frac{(Mx)_i^2}{\sum_j (Mx)_j^2}, \quad i = 0, 1, \dots, d-1.$$

Map of the simplex

$$\frac{x(n+1)}{|x(n+1)|} = Tx(n) \quad \text{and } Tx(n) = Tr_{x(n)} \text{ for } n \geqslant 0.$$

Map of the simplex

$$\frac{x(n+1)}{|x(n+1)|} = Tx(n) \quad \text{and } Tx(n) = Tr_{x(n)} \text{ for } n \geqslant 0.$$

Let $K_{d-1} = \{u \in \mathbb{R}^d : |u| = 1\}$ denote the unit simplex in $\mathbb{R}^d$.

Map of the simplex

$$\frac{x(n+1)}{|x(n+1)|} = Tx(n) \quad \text{and } Tx(n) = Tr_{x(n)} \text{ for } n \geqslant 0.$$

Let $K_{d-1} = \{u \in \mathbb{R}^d : |u| = 1\}$ denote the unit simplex in $\mathbb{R}^d$.

Proposition

The map $T : K_{d-1} \to K_{d-1}$ has a fixed point $u_0 \in K_{d-1}$. If M is primitive, then $u_0 > 0$.

It would be nice if $T : K_{d-1} \to K_{d-1}$ had a unique fixed point u, and if $T^n x \to u$ for any initial nonzero nonnegative vector x, for example $x = (1, 1, \ldots, 1)$.

It would be nice if $T : K_{d-1} \to K_{d-1}$ had a unique fixed point u, and if $T^n x \to u$ for any initial nonzero nonnegative vector x, for example $x = (1, 1, \ldots, 1)$.

The $r(n)$ are useful in the numerical approximation of the tree shift entropy h.

It would be nice if $T : K_{d-1} \to K_{d-1}$ had a unique fixed point u, and if $T^n x \to u$ for any initial nonzero nonnegative vector x, for example $x = (1, 1, \dots, 1)$.

The $r(n)$ are useful in the numerical approximation of the tree shift entropy h.

The entries of the fixed point u might be supposed to carry information about any measure of maximal entropy, but numerical evidence indicates that in general $u(i)$ is *not* the measure of the set of labeled trees that have the symbol i at the root.

In case $k = 2$, $T : K_1 \to K_1$ is a function of a single variable: $T(x, 1 - x) = (1, x^2)/(1 + x^2)$.

In case $k = 2$, $T : K_1 \to K_1$ is a function of a single variable: $T(x, 1-x) = (1, x^2)/(1+x^2)$.

The equations $Tx = x$ reduce to $x^3 + x - 1 = 0$, with root $u = 0.682328$.

In case $k = 2$, $T : K_1 \to K_1$ is a function of a single variable: $T(x, 1-x) = (1, x^2)/(1+x^2)$.

The equations $Tx = x$ reduce to $x^3 + x - 1 = 0$, with root $u = 0.682328$.

The derivative at the fixed point is $T'(u) = -0.635345$, so the fixed point is attracting: $T^n x \to u$ for all $x \in [0, 1]$.

In case $k = 2$, $T : K_1 \to K_1$ is a function of a single variable: $T(x, 1-x) = (1, x^2)/(1+x^2)$.

The equations $Tx = x$ reduce to $x^3 + x - 1 = 0$, with root $u = 0.682328$.

The derivative at the fixed point is $T'(u) = -0.635345$, so the fixed point is attracting: $T^n x \to u$ for all $x \in [0, 1]$.

In the case of the k-tree, alphabet size still $d = 2$, no 11, $T_k : K_1 \to K_1$ is the function $T_k x = 1/(1 + x^k)$.

In case $k = 2$, $T : K_1 \to K_1$ is a function of a single variable: $T(x, 1-x) = (1, x^2)/(1+x^2)$.

The equations $Tx = x$ reduce to $x^3 + x - 1 = 0$, with root $u = 0.682328$.

The derivative at the fixed point is $T'(u) = -0.635345$, so the fixed point is attracting: $T^n x \to u$ for all $x \in [0, 1]$.

In the case of the k-tree, alphabet size still $d = 2$, no 11, $T_k : K_1 \to K_1$ is the function $T_k x = 1/(1 + x^k)$.

There is a critical value $k_0 \approx 4.125$, the solution of

$$k_0 = 1 + k_0^{k_0/(k_0+1)},$$

at which $T'(u) = 1$.

In case $k = 2$, $T : K_1 \to K_1$ is a function of a single variable: $T(x, 1-x) = (1, x^2)/(1+x^2)$.

The equations $Tx = x$ reduce to $x^3 + x - 1 = 0$, with root $u = 0.682328$.

The derivative at the fixed point is $T'(u) = -0.635345$, so the fixed point is attracting: $T^n x \to u$ for all $x \in [0, 1]$.

In the case of the k-tree, alphabet size still $d = 2$, no 11, $T_k : K_1 \to K_1$ is the function $T_k x = 1/(1 + x^k)$.

There is a critical value $k_0 \approx 4.125$, the solution of

$$k_0 = 1 + k_0^{k_0/(k_0+1)},$$

at which $T'(u) = 1$.

We claim that for $k > k_0$ besides the fixed point u there is also an attracting periodic orbit $\{p_1, p_2\}$, and no other periodic points.

T_7, T_7^2

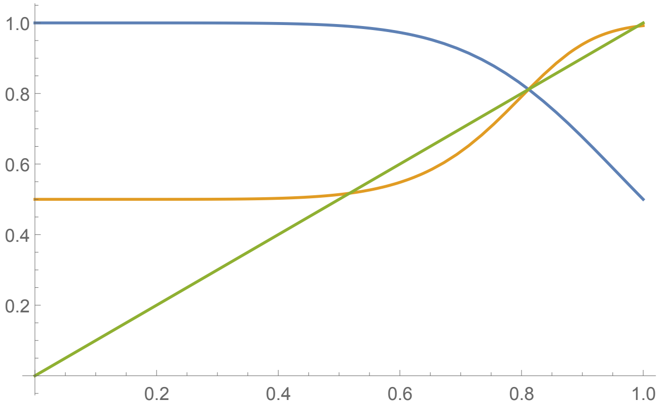

For the 1-dimensional golden mean SFT, the map T on $K_1 = \{(x, y) : x \geqslant 0, y \geqslant 0, x + y = 1\}$ is $T(x, y) = ((x + y)^2, x^2)/((x + y)^2 + x^2)$.

For the 1-dimensional golden mean SFT, the map T on $K_1 = \{(x, y) : x \geqslant 0, y \geqslant 0, x + y = 1\}$ is $T(x, y) = ((x + y)^2, x^2)/((x + y)^2 + x^2)$.

The equations $T(x, y) = (x, y)$ define an algebraic curve.

For the 1-dimensional golden mean SFT, the map T on $K_1 = \{(x,y) : x \geqslant 0, y \geqslant 0, x+y=1\}$ is $T(x,y) = ((x+y)^2, x^2)/((x+y)^2 + x^2)$.

The equations $T(x,y) = (x,y)$ define an algebraic curve.

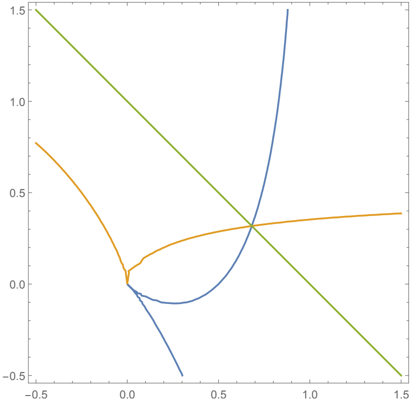

When $d > 3$, looking for the fixed point of T involves intersections of surfaces.

When $d > 3$, looking for the fixed point of T involves intersections of surfaces.

For the system corresponding to the matrix $M = 010, 101, 101$, the plot of the equations for the fixed point of T is

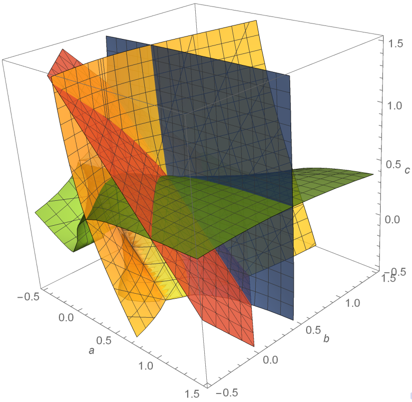

In this case, numerical calculations indicate that iterates of T_3 converge to a fixed point:
$T^n(1,1,1) \to u = (0.188844, 0.405578, 0.405578)$.

In this case, numerical calculations indicate that iterates of T_3 converge to a fixed point:
$T^n(1,1,1) \to u = (0.188844, 0.405578, 0.405578)$.

For the primitive matrix $M = 1100, 0010, 0101, 1000$, with $d = 4$, we find a fixed point
$u = (0.409528, 0.0845909, 0.22457, 0.281312)$,

In this case, numerical calculations indicate that iterates of T_3 converge to a fixed point:
$T^n(1,1,1) \to u = (0.188844, 0.405578, 0.405578)$.

For the primitive matrix $M = 1100, 0010, 0101, 1000$, with $d = 4$, we find a fixed point
$u = (0.409528, 0.0845909, 0.22457, 0.281312)$,

while the iterates $T^n(1,1,1,1)$ oscillate between
$(0.444343, 0.000452842, 0.443437, 0.111767)$ and
$(0.327277, 0.32528, 0.0208321, 0.326611)$.

In this case, numerical calculations indicate that iterates of T_3 converge to a fixed point:
$T^n(1,1,1) \to u = (0.188844, 0.405578, 0.405578)$.

For the primitive matrix $M = 1100, 0010, 0101, 1000$, with $d = 4$, we find a fixed point
$u = (0.409528, 0.0845909, 0.22457, 0.281312)$,

while the iterates $T^n(1,1,1,1)$ oscillate between
$(0.444343, 0.000452842, 0.443437, 0.111767)$ and
$(0.327277, 0.32528, 0.0208321, 0.326611)$.

The description of the dynamics of all these maps T via rigorous analysis presents a serious challenge.