Braids Lecture 26: Subshifts of Finite Type

Our goal in the next few lectures is to establish some strong results about pseudo-Anosovs, based in part on the estimates of the entropy derived previously. But first we look at a "building block" for the pA's — subshifts.

$A = (a_{ij})$ is a $k \times k$ matrix s.t. $a_{ij} = 0$ or 1 ("topological transition matrix") "adjacency matrix", called a 0–1 matrix.

A determines a subshift of finite type as follows.

Let $S_k = \{1, \dots, k\}$, $\Sigma_1(k) = \prod_{i=-\infty}^{\infty} S_k^i$, $S_k^i = S_k$.

S_k is given the discrete topology (every subset an open set), and $\Sigma_1(k)$ the product topology.

The subset $\Sigma_{1A} \subset \Sigma_1(k)$ is the closed subset consisting of of those bi-infinite sequences

$$\underline{b} = (b_n)_{n \in \mathbb{Z}} \quad \text{s.t. } a_{b_i b_{i+1}} = 1, \ \forall i \in \mathbb{Z}.$$

Think of k "states" or boxes:

$\boxed{1}$ $\boxed{2}$ $\boxed{3}$ $\cdots$ $\boxed{k}$

At "time" n, the system can be in one of those states.

$\Sigma_1(k)$ = all possible "histories"

Σ_{1A} = " " " given that the state can move from box i to j iff $a_{ij} = 1$.

In other words, a 3 can only follow a 2 in the sequence if $a_{23} = 1$.

If $A = I$, then Σ_A = sequences consisting of only one number, repeated forever.

The shift $\sigma_A : \Sigma_A \to \Sigma_A$ is defined by

$$\sigma_A[(b_n)_{n\in\mathbb{Z}}] = (b_n')_{n\in\mathbb{Z}}, \text{ with } b_n' = b_{n+1}.$$

[So $b_0' = b_1$, $b_1' = b_2$, etc: shift sequence left.]

σ_A is continuous [The preimage — right shift — of any open set is open.]

Let $C_i = \{\underline{x} \in \Sigma(k) \mid x_0 = i\}$ ← b_i: ∞ sequences with "i" at slot 0.

$D_i = C_i \cap \Sigma_A$ ← same, but restricted to Σ_A.

$\mathcal{D} = \{D_1, \ldots, D_k\}$ is an open cover of Σ_A by pairwise-disjoint elements, since they all differ in their 0^{th} slot. [D_i are open in the relative topology on Σ_A, not necessarily open in $\Sigma(k)$]

Define the norm of $B = (b_{ij})_{1\le i,j\le k}$ by

$$\|B\| = \sum_{i,j=1}^{k} |b_{ij}|.$$

recall part of def'n of topological entropy

Then:

actually the minimum cardinality of a subcover.

$$N_n(\sigma_A, \mathcal{D}) = \mathrm{card}(\mathcal{D} \vee \sigma_A^{-1}\mathcal{D} \vee \cdots \vee \sigma_A^{-n+1}\mathcal{D}) \leq \|A^{n-1}\|$$

In case $D_i = \emptyset$ for some i's.

Check: $N_1(\sigma_A, \mathcal{D}) = \mathrm{card}\,\mathcal{D} \leq k = \|A^0\|$

$N_2(\sigma_A, \mathcal{D}) = \mathrm{card}(\mathcal{D} \vee \sigma_A^{-1}\mathcal{D})$ $\left[\text{Recall } \mathcal{A} \vee \mathcal{B} = \{A_i \cap B_j\}_{\substack{i\in I\\ j\in J}}\right]$

shift right

If we shift an element of $\mathcal{D}$ to the right the number of ways it can have an i in slot 0 depends on the number of nonzero entries in column a_{ij}, $1 \leq j \leq k$. Each of these is an intersection with $\mathcal{D}$, so adds to $\mathrm{card}(\mathcal{D} \vee \sigma_A^{-1}\mathcal{D})$. Hence, $N_2(\sigma_A, \mathcal{D}) \leq \|A\|$.

Similarly, a_{ij}^n counts the number of sequences $(i, i_1, \ldots, i_{n-1}, j)$ which will interact with a given "history" for the past n steps. Here, $a_{i_\ell i_{\ell+1}} = 1$, $1 \leq \ell < n$.

Hence, $$\limsup_{n\to\infty} \frac{1}{n} \log N_n(\sigma_A, \mathcal{D}) \leq \limsup_{n\to\infty} \frac{1}{n} \log \|A^{n-1}\|$$
$$= \limsup_{n\to\infty} \log \|A^n\|^{1/n}$$

This is the log of the spectral radius of A, $\log \lambda$. In fact it is even better than that:

Proposition: For any subshift of finite type $\sigma_A : \Sigma_A \to \Sigma_A$, we have $h(\sigma_A) = \log \lambda$, where λ is the spectral radius of A.

Proof: Each open cover $\mathcal{U}$ of Σ_A is refined by a cover of the form
$$\bigvee_{i=-\ell}^{\ell} \sigma_A^{-i} \mathcal{D}$$
← Just shift the "fixed" value to other positions as needed.

This implies
$$N_{n+1}(\sigma_A, \mathcal{U}) \leq \operatorname{card}\left(\bigvee_{j=-\ell}^{n+\ell} \sigma_A^{-j} \mathcal{D} \right)$$
$$= \operatorname{card}\left(\bigvee_{j=0}^{n+2\ell} \sigma_A^{-j} \mathcal{D} \right)$$
← Shifting all the open sets doesn't change card.
$$= N_{n+2\ell+1}(\sigma_A, \mathcal{D})$$
← By definition of N!

Hence, $h(\sigma_A, \mathcal{U}) \leq h(\sigma_A, \mathcal{D})$.

Since $\mathcal{U}$ is arbitrary, we have $h(\sigma_A) = h(\sigma_A, \mathcal{D})$.

On the previous page we showed $h(\sigma_A, \mathcal{D}) \leq \log \lambda$, with λ the spectral radius of A.

If $D_i \neq \phi$, $1 \leq i \leq k$, this means that each state occurs, and the inequality is an equality, which proves the proposition.

If some states do not occur, we have to show that ultimately $h_A(\sigma_A, \mathcal{D})$ is still dominated by the modulus of the largest eigenvalue of A, which is indeed the case. Since our interest will lie in A's where all states occur, we leave out this part of the proof.

[See FLP p. 192].