PROCEEDINGS OF THE
AMERICAN MATHEMATICAL SOCIETY
Volume 113, Number 4, December 1991

ENTROPY OF SEMIPATTERNS OR HOW TO CONNECT THE DOTS TO MINIMIZE ENTROPY

JONATHAN ASHLEY, ETHAN M. COVEN, AND WILLIAM GELLER

(Communicated by Kenneth R. Meyer)

ABSTRACT. In this paper we solve the following problem: Given a map φ from a finite subset of the reals into the reals, how do you connect the dots in the graph of φ in order to minimize the topological entropy of the resulting map of the interval?

INTRODUCTION

In this paper, we solve the following problem. Given a map φ from a finite subset P of the reals into the reals, how do you connect the dots in the graph of φ in order to minimize the topological entropy of the resulting continuous map of the interval. That is, how do you find an *extension* f_φ of φ (i.e., a continuous map $f_\varphi\colon I \to I$, where $I = [\min(P \cup \varphi(P)),\ \max(P \cup \varphi(P))]$, such that $f_\varphi | P = \varphi$) such that $\operatorname{ent}(f_\varphi) \le \operatorname{ent}(f)$ for every extension f of φ. Here $\operatorname{ent}(\cdot)$ denotes topological entropy.

We use the following terminology, borrowed from [MN]. Let P be a finite subset of the reals. A *pattern* on P is a map from P to itself, a *semipattern* on P is a map from P to the reals. The *entropy* of a semipattern φ on P is $\operatorname{ent}(\varphi) = \inf\{\operatorname{ent}(f)\}$, where the infimum is taken over all extensions f of φ. Thus we are asking how to construct an entropy-minimizing extension of φ.

Throughout this paper, P is a finite subset of the reals, φ a semipattern on P, $\overline{P} = P \cup \varphi(P)$, and $I = [\min \overline{P}, \max \overline{P}]$.

If φ is a pattern on P, then the solution is well known: $\operatorname{ent}(\varphi)$ is the logarithm of the largest real eigenvalue of the "adjacency" matrix A of φ, defined as follows. If $P = \{p_1 < \cdots < p_k\}$, then $A = (a_{ij})$ is $(k-1) \times (k-1)$ and $a_{ij} = 1$ if p_j and p_{j+1} are in the closed interval whose endpoints are $\varphi(p_i)$ and $\varphi(p_{i+1})$, $a_{ij} = 0$ otherwise. Furthermore $\operatorname{ent}(\varphi) = \operatorname{ent}(f)$ for every P-monotone extension f of φ. Here $f\colon I \to I$ is *P-monotone* if P contains the endpoints of I and f is (not necessarily strictly) monotone on intervals whose endpoints are consecutive members of P. Analogously, $f\colon I \to I$ is

Received by the editors April 8, 1990; presented to the 851st meeting of the American Mathematical Society on October 22, 1989.

1980 *Mathematics Subject Classification* (1985 *Revision*). Primary 58F08, 54H20.

©1991 American Mathematical Society
0002-9939/91 \$1.00 + \$.25 per page

P-linear if P contains the endpoints of I and f is affine on intervals whose endpoints are consecutive members of P. There is only one P-linear extension of φ, which we denote by L_φ.

If φ is a pattern on P, then the "obvious" solution—connect the dots with straight lines, i.e., $f_\varphi = L_\varphi$—is correct. The following example shows that, in general, the "obvious" solution may be wrong. Let $P = \{0, \frac{1}{2}, 1\}$, $\varphi(0) = \varphi(1) = 0$, and $\varphi(\frac{1}{2}) = a$, where $\frac{1}{2} < a < 1$. Consider the following two maps.

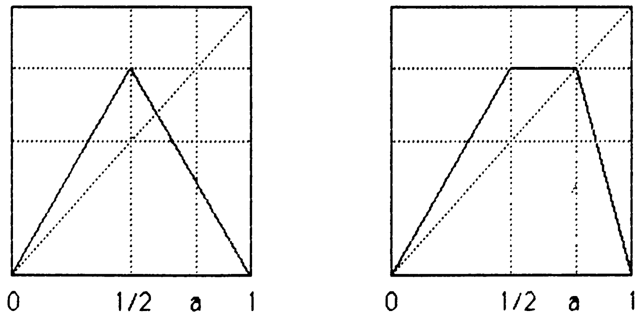

The map on the left is the P-linear extension of φ. The map on the right is the $\overline{P}$-linear extension of the pattern $\overline{\varphi}$ on $\overline{P} = \{0, \frac{1}{2}, a, 1\}$ defined by $\overline{\varphi}(0) = \overline{\varphi}(1) = 0$ and $\overline{\varphi}(\frac{1}{2}) = \overline{\varphi}(a) = a$. Using the Misiurewicz–Szlenk formula [MS] for the entropy of a piecewise monotone map,

$$\operatorname{ent}(f) = \lim_{n\to\infty} (1/n) \log \operatorname{var}(f^n),$$

the map on the left has entropy $\log 2a$. The largest real eigenvalue of the adjacency matrix of $\overline{\varphi}$ is one, so the map on the right has entropy zero. We show that this case is typical in that:

Theorem. *Let φ be a semipattern on P. Then there is a pattern $\overline{\varphi}$ on $\overline{P}$ such that $\overline{\varphi}|P = \varphi$ and $\operatorname{ent}(L_{\overline{\varphi}}) \le \operatorname{ent}(f)$ for every extension f of φ. Furthermore, there is an algorithm for constructing such a pattern and hence for constructing an entropy-minimizing map.*

1. Preliminary results

Lemma 1. *Let $f, g\colon [a, b] \to [a, b]$ be continuous. If $f = g$ except on an at most countable collection of intervals, on each of which g is constant, then* $\operatorname{ent}(g) \le \operatorname{ent}(f)$.

Proof. It suffices to prove the lemma in the case that the collection consists of a single interval. Then it holds when the collection consists of finitely many intervals. Now suppose that the collection of intervals is $\{J_1, J_2, \dots\}$. Let $g_n = g$ on $J_1 \cup \dots \cup J_n$ and $g_n = f$ off $J_1 \cup \dots \cup J_n$. Then $g_n \to g$ in the C^0-topology and $\operatorname{ent}(g_n) \le \operatorname{ent}(f)$ for every n. Since $\operatorname{ent}(\cdot)$ is lower semicontinuous [M1, Theorem 2], $\operatorname{ent}(g) \le \operatorname{ent}(f)$.

So let g be constant on J and $g = f$ off J. Recall Bowen's formula for entropy: $\text{ent}(f) = \lim_{\varepsilon\to 0} \limsup[(1/n)\log s_f(n, \varepsilon)]$, where $s_f(n, \varepsilon)$ is the largest cardinality of an (n, ε)-separated set for f. Here S is (n, ε)-*separated* for f if for every $x \neq y$ in S, $|f^i(x)-f^i(y)| > \varepsilon$ for some i, $0 \le i \le n-1$. To show that $\text{ent}(g) \le \text{ent}(f)$, it suffices to show that $\limsup[(1/n)\log s_g(n, \varepsilon)] \le \limsup[(1/n)\log s_f(n, \varepsilon)]$ for every $\varepsilon > 0$.

Let $n \ge 2$ and let S be an (n, ε)-separated set for g. For $k = 0, 1, \dots, n-2$, let $E_k = \{x \in S\colon k$ is the least integer such that $f^k(x) \in J\}$ and let $E_{n-1} = S - \bigcup_{k=0}^{n-2} E_k$. Then E_k is $(k+1, \varepsilon)$-separated for f $(k = 0, \dots, n-1)$. It follows that $s_g(n, \varepsilon) \le s_f(1, \varepsilon) + \cdots + s_f(n, \varepsilon)$. The result then follows from the elementary fact that if $a_n, b_n > 0$ and $b_n \le a_1 + \cdots + a_n$ for all $n \ge 1$, then $\limsup[(1/n)\log b_n] \le \limsup[(1/n)\log a_n]$. □

Lemma 2. $\text{ent}(\varphi) = \inf\{\text{ent}(f) | f$ *is a* P-*monotone extension of* $\varphi\}$.

Proof. We first show that the infimum in the definition of $\text{ent}(\varphi)$ is unchanged if we require that P contains the endpoints of I. Let $Q = P \cup \{a, b\}$, where $I = [a, b]$, and define a semipattern ψ on Q by $\psi = \varphi$ on P, $\psi(a) = \varphi(p_{\min})$, and $\psi(b) = \varphi(p_{\max})$. Here $p_{\min}$ and $p_{\max}$ denote the smallest and largest members of P. Suppose that f is an extension of φ. Define $g\colon I \to I$ by

$$g(x) = \begin{cases} f(p_{\min}) & \text{if } x \in [a, p_{\min}], \\ f(x) & \text{if } x \in [p_{\min}, p_{\max}], \\ f(p_{\max}) & \text{if } x \in [p_{\max}, b]. \end{cases}$$

Then g is a Q-monotone extension of ψ and by Lemma 1, $\text{ent}(g) \le \text{ent}(f)$.

Now suppose that f is an extension of φ and P contains the endpoints of I. Let $p < p'$ be consecutive members of P. Define $g\colon I \to I$ as follows. Let $g = f$ off $[p, p']$. If $f(p) = f(p')$, let $g(x) = f(p)$ for $x \in [p, p']$. If $f(p) < f(p')$, let

$$g(x) = \begin{cases} f(p) & \text{if } x \in [p, q], \\ \sup\{f(y) | q \le y \le x\} & \text{if } x \in [q, r], \\ f(p') & \text{if } x \in [r, p'], \end{cases}$$

where $q = \sup\{x \in [p, p'] | f(x) = f(p)\}$ and $r = \inf\{x \in [q, p'] | f(x) = f(p')\}$. If $f(p) > f(p')$, replace "$q \le y \le x$" by "$x \le y \le r$" in the definition of g. The graph of g on $[q, r]$ is obtained by "pouring water" into the graph of f on $[q, r]$ (cf. [ALMS]).

Repeating this procedure for each such pair $p < p'$ yields a P-monotone extension g of φ. By Lemma 1, $\text{ent}(g) \le \text{ent}(f)$. □

An extension f of φ is P-*Markov* if it is P-monotone and every member of P has a finite orbit. If f is a P-Markov extension of φ, then there is an integer N such that $Q = P \cup f(P) \cup \cdots \cup f^N(P)$ is f-invariant. In this case, $f|Q$ is a pattern on Q.

Lemma 3. *If f is P-monotone extension of φ, then there are P-Markov extensions g of φ with* $\mathrm{ent}(g)$ *arbitrarily close to* $\mathrm{ent}(f)$. *Therefore* $\mathrm{ent}(\varphi) = \inf\{\mathrm{ent}(f) | f$ *is a P-Markov extension of* $\varphi\}$.

To prove Lemma 3, we use some results from kneading theory. Following [MT], let $\mathscr{C}_\sigma(T)$ denote the set of continuous, piecewise monotone self-maps of I with turning point set T, alternating strictly increasing and strictly decreasing on the laps, and increasing or decreasing on the first lap accordingly as $\sigma = +1$ or -1.

Lemma 4 (cf. [MT, Lemma 12.3]). *Let* $f \in \mathscr{C}_\sigma(T)$. *For each turning point* c, *define*

$N(c) =$ *number of points in the orbit of* c *if* c *has a finite orbit,*

$N(c) =$ *largest* n *such that* $f^n(c)$ *is a turning point if* c *has an infinite orbit,*

and let $N = \max\{N(c)\}$. *Then* $\mathrm{ent}(g)$ *tends to* $\mathrm{ent}(f)$ *as* g *tends to* f *in the* C^0*-topology, keeping* $g \in \mathscr{C}_\sigma(T)$ *and* $g = f$ *on* $T \cup f(T) \cup \cdots \cup f^N(T)$.

Proof. The condition that $g = f$ on $T \cup f(T) \cup \cdots \cup f^N(T)$ serves the same function as the hypothesis in the proof of [MT, Lemma 12.3] that no turning point is periodic; namely, to guarantee that the coefficients of the kneading determinants of f and g can be made to agree on an arbitrarily long initial segment. The proof is then completed by the observation that entropy depends continuously on the kneading determinant. □

The published kneading theory (e.g., [MT]) is stated for maps that are piecewise strictly monotone. However, as M. Misiurewicz has pointed out, those parts of the kneading theory that do not use smoothness work for maps that are not necessarily strictly monotone on the pieces [M2, p. 225]. One merely replaces turning points by "turning intervals" (which, of course, may be degenerate). For the correct formalism, see [M2].

Let $\mathscr{C}'_\sigma(T)$ denote the set of continuous, piecewise (but not necessarily strictly) monotone self-maps of I with T the collection of turning intervals, alternating increasing and decreasing on the laps (i.e., the components of the complement of the union T^* of the turning intervals), and increasing or decreasing on the first lap according as $\sigma = +1$ or -1. In this setting, Lemma 4 becomes

Lemma 5. *Let* $f \in \mathscr{C}'_\sigma(T)$. *For each turning interval* C, *define*

$N(C) = 1+$*number of points in the orbit of* $f(C)$ *if* $f(C)$ *has a finite orbit,*

$N(C) =$ *largest* n *such that* $f^n(C)$ *is in a turning interval if* $f(C)$ *has an infinite orbit,*

and let $N = \max\{N(C)\}$. *Then* $\mathrm{ent}(g)$ *tends to* $\mathrm{ent}(f)$ *as* g *tends to* f *in the* C^0*-topology, keeping* $g \in \mathscr{C}'_\sigma(T)$ *and* $g = f$ *on* $T^* \cup f(T^*) \cup \cdots \cup f^N(T^*)$.

Proof of Lemma 3. It suffices to show that if $f \in \mathscr{C}'_\sigma(T)$ is P-monotone, then arbitrarily C^0-close to f there are P-Markov maps $g \in \mathscr{C}'_\sigma(T)$ which agree with f on $T^* \cup f(T^*) \cup \cdots \cup f^N(T^*)$, where N is as in Lemma 5.

Let $f \in \mathscr{C}'_\sigma(T)$ be P-monotone and let $p \in P$ have an infinite orbit. There exist m, n and an open interval $(a, b) \subseteq I$ containing $f^m(p)$ and $f^n(p)$ such that

(1) $N < m < n$,
(2) $b - a$, $|f^m(p) - f^n(p)|$, and $|f^{m+1}(p) - f^{n+1}(p)|$ are small, and
(3) $(a, b) \cap [T^* \cup f(T^*) \cup \cdots \cup f^N(T^*)] = \{f^m(p), f^n(p)\}$.

Define f_1 as follows: $f_1 = f$ off (a, b), $f_1(x) = f(x)$ for $x = a, b, f^m(p)$, $f_1(f^n(p)) = f(f^m(p))$, and f_1 is linear on the intervals whose endpoints are consecutive members of $\{a, b, f^m(p), f^n(p)\}$. Then $f_1 \in \mathscr{C}'_\sigma(T)$ is P-monotone, agrees with f on $T^* \cup f(T^*) \cup \cdots \cup f^N(T^*)$, and the number of points in P having an infinite f_1-orbit is at least one less than the corresponding number for f.

Repeat this argument with f_1 in place of f and $N_1 = n$ in place of N. A finite number of iterations of this procedure yields a P-Markov map $g \in \mathscr{C}'_\sigma(T)$, which agrees with f on $T^* \cup f(T^*) \cup \cdots \cup f^N(T^*)$. By choosing appropriate meanings for the various occurrences of the word "small," we can make g as C^0-close to f as we wish. □

2. The theorem

For a nonnegative matrix A, let $\lambda(A)$ denote the largest real eigenvalue of A. We need the following matrix-theoretic result, which may be of independent interest.

Lemma 6. *Let A be a nonnegative matrix and S a subset of the indexing set of A. Then there exists $k \in S$ such that if B is defined by*

$$b_{ij} = a_{ij} \quad \text{if } i \notin S,$$
$$b_{ij} = 0 \quad \text{if } i \in S,\ i \neq k,$$
$$b_{kj} = \sum_{i \in S} a_{ij},$$

then $\lambda(B) \le \lambda(A)$.

Proof. For $n \ge 1$, let $A_n = A + (1/n)J$, where J is any positive matrix. Since A_n is positive, it has a positive left eigenvector $l^{(n)}$ corresponding to $\lambda(A_n)$. There exists $k_n \in S$ such that $l^{(n)}_{k_n} \le l^{(n)}_i$ for all $i \in S$. By passing to a subsequence, we may assume that k_n is the same for all n. Let k be this common value and let B be as in the statement of the lemma.

Since $A_n \to A$, it follows that $\lambda(A_n) \to \lambda(A)$. Thus it suffices to show that $\lambda(B) \le \lambda(A_n)$ $(n \ge 1)$. It is easy to verify that $l^{(n)}B \le l^{(n)}A_n$. Let r be a nonnegative right eigenvector for B corresponding to $\lambda(B)$. Then $\lambda(B)l^{(n)}r = l^{(n)}Br \le l^{(n)}A_n r = \lambda(A_n)l^{(n)}r$. Since $l^{(n)}r > 0$, $\lambda(B) \le \lambda(A_n)$. □

Theorem. *Let* φ *be a semipattern on* P. *Then there is a pattern* $\overline{\varphi}$ *on* $\overline{P}$ *such that* $\overline{\varphi}|P = \varphi$ *and* $\mathrm{ent}(L_{\overline{\varphi}}) \leq \mathrm{ent}(f)$ *for every extension* f *of* φ. *Furthermore, there is an algorithm for constructing such a pattern and hence for constructing an entropy-minimizing map.*

Proof. We first show that if f is a P-Markov extension of φ, then there is a pattern ψ on $\overline{P}$ such that $\psi|P = \varphi$ and $\mathrm{ent}(L_\psi) \leq \mathrm{ent}(f)$. The first conclusion of the theorem then follows from Lemma 3 and the fact that there are only finitely many patterns on $\overline{P}$.

Let f be a P-Markov extension of φ and let $Q = P \cup f(P) \cup \cdots \cup f^N(P)$ be f-invariant. Then $\mathrm{ent}(f) = \log \lambda(A)$, where A is the adjacency matrix A of the pattern $f|Q$.

Write $Q = \{q_1 < \cdots < q_m\}$. Let $\{(i_k, j_k): k = 1, \dots, r\}$ be a list of all pairs $i < j$ such that $q_i, q_j \in P$ but $q_{i+1}, \dots, q_{j-1} \notin P$. Let B_1 be the matrix obtained by applying Lemma 6 to A and $S = \{q_{i_1}, q_{i_1} + 1, \dots, q_{j_1 - 1}\}$. Let B_2 be the matrix obtained by applying Lemma 6 to B_1 and $S = \{q_{i_2}, q_{i_2} + 1, \dots, q_{j_2 - 1}\}$. Continuing this way, we obtain a matrix B_r. It follows from the way that B_r was obtained that it is the adjacency matrix of some pattern θ on Q. Thus $\mathrm{ent}(L_\theta) \leq \mathrm{ent}(f)$. Since $\theta(Q) \subseteq \overline{P}$, $\theta|\overline{P}$ is a pattern, which we denote by ψ. Clearly $\psi|P = \varphi$. L_θ is a Q-monotone extension of θ and hence is a $\overline{P}$-monotone extension of ψ. Therefore $\mathrm{ent}(L_\theta) = \mathrm{ent}(L_\psi)$ and so $\mathrm{ent}(L_\psi) \leq \mathrm{ent}(f)$.

We now turn to the algorithm for finding $\overline{\varphi}$. Let $\Psi = \{\psi: \psi$ is a pattern on $\overline{P}$, $\psi|P = \varphi\}$. For $\psi \in \Psi$, let A_ψ be the adjacency matrix of ψ. Finally, let $\mathscr{A} = \{A_\psi: \psi \in \Psi\}$.

Sturm's Theorem [vdW, p. 220] gives an algorithm for determining the number of distinct real roots of a polynomial with integer coefficients in a closed interval with rational endpoints. From this it is straightforward to construct an algorithm, which, given two polynomials with integer coefficients and having real roots, determines whether their largest real roots are equal or not, and if not, which of the two is smaller. Thus there is an algorithm for determining which among a finite set of nonnegative matrices with integer entries has the smallest largest real eigenvalue. Applying this algorithm to $\mathscr{A}$ yields a pattern $\psi \in \Psi$ such that $\lambda(A_\psi) \leq \lambda(A_{\psi'})$ for every $\psi' \in \Psi$. Let $\overline{\varphi}$ be any such ψ. □

3. Implementing the algorithm: an example

In this section we give an example to show that we need only look at a relatively small proportion of the patterns on $\overline{P}$ to find one whose $\overline{P}$-linear extension is an entropy-minimizing map in the Theorem. The interested reader may supply the formal complexity-theoretic arguments himself for the general case.

Example. Let $P = \{1, 2, 4, 5, 8, 9\}$, $\varphi(1) = 9$, $\varphi(2) = 6$, $\varphi(4) = 1$, $\varphi(5) = 7$, $\varphi(8) = 3$, $\varphi(9) = 9$. Here $\overline{P} = \{1, \dots, 9\}$ and $\overline{P} - P = \{3, 6, 7\}$. Thus there are $9^3 = 729$ patterns $\overline{\varphi} = \overline{P}$ such that $\overline{\varphi}|P = \varphi$. Looking only at those whose $\overline{P}$-linear extensions are P-monotone (Lemma 2), there are 6 choices for $\overline{\varphi}(3)$ and $5+4+3+2+1 = 15$ choices for $(\overline{\varphi}(6), \overline{\varphi}(7))$, giving 90 such patterns. On the other hand, if $\overline{p} \in \overline{P} - P$ and p, p' are the consecutive members of P such that $p < \overline{p} < p'$, then, as in the proof of the Theorem, we may choose $\overline{\varphi}$ so that $\overline{\varphi}(\overline{p}) = \varphi(p)$ or $\varphi(p')$. There are only two such choices for $\overline{\varphi}(3)$ and only three such choices for $(\overline{\varphi}(6), \overline{\varphi}(7))$, giving just $2 \cdot 3 = 6$ patterns to look at.

Acknowledgments

The authors benefited from discussions with a number of people, especially with K. Brucks concerning kneading theory; with Z. Nitecki, who first pointed out to us that connecting the dots with straight lines need not minimize entropy; and with the members of Project 4064 at the IBM Almaden Research Center concerning implementing the algorithm. Much of this work was done during the Workshop on Ergodic Theory and Symbolic Dynamics held at the University of Washington in summer 1989. The second and third authors are grateful to the organizers for invitations and support.

References

[ALMS] L. Alseda, J. Llibre, M. Misiurewicz, and C. Simo, *Rotation intervals for a class of maps of the real line to itself*, Ergodic Theory Dynamical Systems **6** (1986), 117–132.

[MT] J. Milnor and W. Thurston, *On iterated maps of the interval*, Lecture Notes in Math., vol. 1342, Springer-Verlag, New York, pp. 465–563.

[M1] M. Misiurewicz, *Horseshoes for mappings of the interval*, Bull. Acad. Polon. Sci. Ser. Sci. Math. **27** (1979), 167–169.

[M2] ——, *Jumps of entropy in one dimension*, Fund. Math. **132** (1989), 215–226.

[MN] M. Misiurewicz and Z. Nitecki, *Combinatorial patterns for maps of the interval*, Mem. Amer. Math. Soc. (to appear).

[MS] M. Misiurewicz and W. Szlenk, *Entropy of piecewise monotone mappings*, Studia Math. **67** (1980), 45–63.

[vdW] B. van der Waerden, *Modern algebra* vol. I, Frederick Ungar, New York, 1953.

Department of Mathematics, University of Washington, Seattle, Washington 98195

Current address: IBM Almaden Research Center, San Jose, California 95120

Department of Mathematics, Wesleyan University, Middletown, Connecticut 06459

Institute of Mathematics, The Hebrew University of Jerusalem, Jerusalem, Israel

Current address: Department of Mathematics, University of Maryland, College Park, Maryland 20742