

Computable Ergodic Optimisation

Gayral Léo, Hoyrup Mathieu

Abstract

Links between physical systems and computability properties have been an active field of investigation in recent years. Inspired by a previous work [GST25] in the context of positive temperature Gibbs measures, we prove here that in the context of zero-temperature ergodic optimisation, for a computable potential and provided with several reasonable assumptions, the maximum ergodic average is a computable real number, and the set of maximising measures is a Π_1 -computable compact set.

Then, in the more specific context of symbolic dynamics, with finite-range interactions on subshifts of finite type, we provide an explicit algorithm to compute both the maximum ergodic average and the set of maximising measures in finite time, with a matching code repository [Gay26].

1 Introduction

Ergodic optimisation can be interpreted as the study of the zero-temperature limit behaviours of a thermodynamic model, in statistical physics. These theoretical arguments help us reach a better understanding of what actually happens in a physical material at low temperatures, notably regarding the emergence of (quasi)crystalline (a)periodic structures. We redirect interested readers to Jenkinson’s relatively recent survey [Jen19] for more details and context on many typical results on ergodic optimisation, as we will here focus on a quite barebones general definition, to better highlight the seldom studied computational aspects.

We start from a *dynamical system* (X, d, T) , where (X, d) is a compact metric space, and $T : X \rightarrow X$ a continuous transformation. Let $\mathcal{M}_T$ the set of *T-invariant* probability measures (*i.e.* such that $\mu \circ T^{-1} = \mu$), compact in the weak-* topology. Given a continuous *potential* function $\varphi : X \rightarrow \mathbb{R}$, we can thus define the *maximum ergodic average* $\beta(\varphi) = \max_{\mu \in \mathcal{M}_T} \int \varphi d\mu$ among *T*-invariant probability measures, and the (convex and compact) set of such *φ -maximising measures* $\mathcal{M}_{\max}(\varphi) = \{\mu \in \mathcal{M}_T : \int \varphi d\mu = \beta(\varphi)\}$.

Computable analysis aims at providing a general framework and tools to computationally describe analytic objects from mathematics, such as real numbers or closed subsets for instance. For general considerations on computable analysis, we refer the readers to the classical textbook by Weihrauch [Wei00], as well as more recent works for more details on links with dynamical systems [Col20] and measure and probability theory [HR21]. Computability theory has permitted the characterisation of the computably realisable values of many conjugacy invariants [HM10; Mey11; JV15; Boy+15; Wes17; HS18; ENT23] of dynamical systems, providing practical tools to distinguish non-equivalent explicitly defined systems. Beyond these invariants, broader links between computability and dynamical systems represent an active area of research [BGZ11; BW24; CR24; GZ24].

A notable fact, that will represent an unavoidable limitation for broad general results about ergodic optimisation, is that there exist some “simple” computable dynamical systems for which the invariant measures are known to be *non-computable* [GHR11; BGR12] (and that’s before taking into account any further information regarding the potential, for ergodic optimisation).

A few existing results concern more specifically the interactions between statistical physics and computability, such as the computability of $\beta \mapsto \sup_{\mu \in \mathcal{M}_T} h(\mu) - \beta \int \varphi d\mu$ (the *pressure* function) on some classes of one-dimensional subshifts [Spa08; Bur+22] and higher-dimensional full shifts [GST25], or of some thermodynamic invariants like the residual entropy [BW20]. The general framework for computable analysis, and the specific case of computable dynamical systems will be properly introduced in Section 2.

In the context of ergodic optimisation, the space of potentials we consider (continuous, Lipschitz, Hölder, finite-range...) has a huge influence on the typical (*i.e.* generic) behaviour of the model. For instance, in the general case of continuous potentials, we have many properties that hold generically (uniqueness of the maximising measure, which has full support, zero entropy, weak but *not* strong mixing...) [Jen06; Bré08; Mor10; EM20]. In comparison, if the map T is *expanding* (such as the shift on a one-sided space $\mathcal{A}^{\mathbb{N}}$, or the multiplication $\times 2$ on the circle $\mathbb{R}/\mathbb{Z}$), generically for Lipschitz potentials, the maximising measures have a *periodic* support instead [Con16].

In practice, a lot of the models we consider, those we define in an explicit way, prove to be computable. Moreover, regardless, it is a prerequisite for us to be able to do any kind of numeric simulation. Hence, it makes sense to restrict ourselves to the class of *computable potentials* (which is notably *not* a real vector space, as not all real numbers are computable) and look at what computable properties follow. We might thus assume that φ (*i.e.* the “input” of the problem) is computable, and study the complexity of $\beta(\varphi)$ and $\mathcal{M}_{\max}(\varphi)$ (*i.e.* the “outputs” of the problem). Using mostly folkloric computability tools, in Section 3, we identify several relevant properties for a *computable* dynamical system such that the maximum ergodic average $\beta(\varphi)$ is computable, with a Π_1 upper bound on the complexity of $\mathcal{M}_{\max}(\varphi)$. These relevant properties are then studied on a few representative examples in Section 4.

Even though we have guarantee β is computable under these hypotheses, the corresponding algorithm is unsatisfying, as it informally requires to enumerate $\mathbb{N}$ and wait until *something* happens (at an unknown point in time, simply known to be finite). Hence, in Section 5, in the more specific case of symbolic dynamics, on a subshift of finite type $X_{\mathcal{F}} \subset \mathcal{A}^{\mathbb{N}}$ (explicitly defined by a set of forbidden patterns $\mathcal{F}$), with a finite-range potential $\varphi : \mathcal{A}^r \rightarrow \mathbb{R}$, we explicit an algorithm computing new forbidden patterns $\mathcal{F}(\varphi) \supset \mathcal{F}$ such that $\mathcal{M}_{\max}(\varphi) = \mathcal{M}_T(X_{\mathcal{F}(\varphi)})$.

2 Background on Computability

In this section, we introduce from scratch computability theory while trying to stay as far from the technical details as reasonably possible, in order to define all the necessary notions for our results. Most of the notions in this section are well-documented in textbooks [Wei00; BP03], so we will only provide references for more specific results.

2.1 Computability on the Natural Numbers

There are many equivalent notions of computability (Turing machines, lambda calculus, general recursive functions...). For the sake of simplicity, we will here invoke a very high level view of the notion, with computable functions seen as those represented by an *algorithm* (using the usual *if/else* statements and *for/while* loops), provided with a natural integer input, and expecting a natural integer output.

Notably, a computable function $f : \text{dom}(f) \rightarrow \mathbb{N}$ can be partial, with $x \in \text{dom}(f)$ if and only if the corresponding algorithm *halts* when given $x \in \mathbb{N}$ as an input, in which case it outputs $f(x) \in \mathbb{N}$. We say that f is total when $\text{dom}(f) = \mathbb{N}$.

We fix an explicit bijection $\langle \cdot \rangle : \mathbb{N}^* \rightarrow \mathbb{N}$, that encodes any tuple as an integer. This bijection (and other explicit ones such as $\mathbb{N} \cong \mathbb{Z}$ or $\mathbb{N} \cong \mathbb{Q}$) allows us to transpose all the computability notions from $\mathbb{N}$ to $\mathbb{N}^*$ and more broadly to any reasonable countable set.

A set $X \subset \mathbb{N}$ is said to be **computably enumerable (c.e.)** if it is the image set $f(\mathbb{N})$ of a computable (partial) function f . Equivalently, X is c.e. if and only if it is the domain of a computable function. A set $X \subseteq \mathbb{N}$ is **computable** if its characteristic function $\mathbb{1}_X : \mathbb{N} \rightarrow \{0, 1\}$ is computable *and total*. Equivalently, X is computable if and only if both X and its complement $\mathbb{N} \setminus X$ are c.e. One of the fundamental results from computability theory is the existence of c.e. sets that are not computable (this fact was already palpable in Turing's seminal work [Tur37], though the question of who actually proved/stated what is much more nuanced [HN26]).

We can broaden these notions into a hierarchy of decision problems, with $X \subset \mathbb{N}$ being a Σ_k^0 set if it satisfies $x \in X \Leftrightarrow \exists n_1 \in \mathbb{N}, \forall n_2 \in \mathbb{N}, \exists n_3 \in \mathbb{N}, \dots, (x, n_1, \dots, n_k) \in A$ with k alternating $\exists \forall$ quantifiers over $\mathbb{N}$ and $A \subseteq \mathbb{N}^{k+1}$ a computable set. The c.e. sets are precisely the Σ_1^0 sets.

Complementarily, X is Π_k^0 if the leftmost quantifier is a $\forall$ instead of a $\exists$ (or, equivalently, if the complementary set $\mathbb{N} \setminus X$ is Σ_k^0).

2.2 Computability on the Real Numbers

Now, most topological spaces used in analysis, such as real numbers, are *not* countable. Hence, we need to broaden the previous notions of standard computability in order to properly do computable analysis.

Definition 2.1 (Hierarchy of Computable Real Numbers). *A real number is Σ_1 -computable if it is the supremum of a computable sequence of rational numbers (i.e. a total computable function $f : \mathbb{N} \rightarrow \mathbb{Q}$). Likewise, a real number is Π_1 -computable if it is the infimum of such a sequence.*

*A Δ_1 -computable real number x (simply called **computable** from now on) is both Σ_1 and Π_1 . Equivalently, there is a total computable function $f : \mathbb{N} \rightarrow \mathbb{Q}$ that gives arbitrarily good approximations of x , such that $|x - f(n)| \leq \frac{1}{2^n}$.*

We can extend these notions into a broader hierarchy, where $x \in \mathbb{R}$ is Σ_k if there is a computable $f : \mathbb{N}^k \rightarrow \mathbb{Q}$ such that $x = \sup_{n_1 \in \mathbb{N}} \inf_{n_2 \in \mathbb{N}} \dots f(n_1, n_2, \dots, n_k)$ with k alternating sup and inf (resp. Π_k if the leftmost one is inf).

This particular notation is not standard in computable analysis, where Σ_1 real numbers are usually called left-computably enumerable or lower semi-computable. Still, the notation we use appears in a few articles and has the advantage of reflecting a close relationship between the hierarchy of real numbers and the hierarchy we obtained before for subsets of $\mathbb{N}$, with sup instead

of $\forall$ and $\inf$ instead of $\exists$. This correspondence is actually explicit, through the definition of real numbers as rational cuts [Wei00, Lemma 4.1.18]:

Proposition 2.2. *A real number x is Σ_k -computable if and only if the set $\{q \in \mathbb{Q} : q < x\} \subset \mathbb{Q}$ is Σ_k^0 . A real number x is Π_k if and only if the set $\{q \in \mathbb{Q} : q > x\} \subset \mathbb{Q}$ is Σ_k^0 .*

Broadly speaking, if a computability notion has been defined for objects in a set X , witnessed by the existence of a computable function $f : \mathbb{N} \rightarrow \mathbb{N}$, then a countable family $(x_i)_{i \in \mathbb{N}}$ of such objects is said to be **uniformly computable** if there is a single computable function $f : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that for each $i \in \mathbb{N}$, $f(i, \cdot) : \mathbb{N} \rightarrow \mathbb{N}$ witnesses the computability of x_i . It is stronger than requiring each x_i to be computable individually.

For example, any real number $x \in [0, 1]$ can be expressed as $x = \sum_{i \in \mathbb{N}} \frac{b_i}{2^i}$ with each $b_i \in \{0, 1\}$ being trivially computable in $\mathbb{N}$, but the sequence $b : \mathbb{N} \rightarrow \{0, 1\}$ being uniformly computable directly implies x itself is a computable real number (which is not the case for most real numbers, for all but a countable amount of them).

2.3 Computable Metric Spaces

We now explain how more general abstract spaces can be endowed with a computability structure.

Definition 2.3. *A **computable metric space** is a triplet (X, d, S) where (X, d) is a separable metric space and $S = \{s_i : i \in \mathbb{N}\}$ is a dense subset indexed by $\mathbb{N}$, such that the real numbers $(d(s_i, s_j))_{i, j \in \mathbb{N}}$ are uniformly computable.*

To be precise, the uniformly computable condition means that there is a total computable function $\delta : \mathbb{N}^3 \rightarrow \mathbb{Q}^+$ such that $\delta(i, j, n)$ is a $\frac{1}{2^n}$ -approximation of $d(s_i, s_j)$.

In a computable metric space (X, d, S) , a point $x \in X$ is **computable** if there exists a computable sequence $(i_n)_{n \in \mathbb{N}}$ such that $d(s_{i_n}, x) < 2^{-n}$ for all n . For instance, in $(\mathbb{R}, d, \mathbb{Q})$ where d is the Euclidean metric and $\mathbb{Q}$ comes with a natural indexing, a real number is computable in this sense if and only if it is computable in the sense of Definition 2.1.

In such a space, a **rational ball** is an open ball $B(s_i, r) = \{x \in X : d(s_i, x) < r\}$ where r is a positive rational number. From a computable enumeration $(r_j)_{j \in \mathbb{N}}$ of the positive rational numbers, we derive a computable enumeration $(B_k)_{k \in \mathbb{N}}$ of the rational balls: $B_k = B(s_i, r_j)$ where $k = \langle i, j \rangle$. We say that $U \subseteq X$ is an **effective open set** if there exists a c.e. set $E \subseteq \mathbb{N}$ such that $U = \bigcup_{k \in E} B_k$. Computability of points can be reformulated using rational balls as follows: a point $x \in X$ is computable if and only if the set $\{k \in \mathbb{N} : x \in B_k\}$ is c.e. (this alternative definition can be used to define computable points in topological spaces that are not metrizable, but still countably-based).

An analog of the computability hierarchy of subsets of $\mathbb{N}$ can be defined for closed subsets of a computable metric space.

Definition 2.4. *Let (X, d, S) be a computable metric space and $k \geq 1$. A closed set $F \subseteq X$ is Π_k if there exists a Σ_k^0 set $E \subseteq \mathbb{N}$ such that $X \setminus F = \bigcup_{i \in E} B_i$.*

In particular, a closed set is Π_1 if its complement is an effective open set. One could similarly define the Σ_k closed sets, but it turns out that they are much less useful.

Remark 2.5. The notation Π_k for closed sets should not be confused with the standard notion of general Π_k^0 sets from effective descriptive set theory, which are defined from the effective open sets and using alternating countable intersections and unions rather than $\forall\exists$ quantifiers (we do not use these notions here, so we will not define them properly).

It is true that every Π_k set is Π_k^0 , but conversely Π_k^0 sets are not necessarily closed, and even then they are not Π_k in general (unless $k = 1$).

In the standard terminology from computable analysis, a Π_k set is a set that is Π_1^0 with oracle $\emptyset^{(k-1)}$, where $\emptyset^{(0)}$ is the empty set and $\emptyset^{(n+1)}$ is inductively defined as the halting problem for Turing machines having an access to $\emptyset^{(n)}$ as an oracle. Similarly, as mentioned earlier the Σ_k and Π_k real numbers are usually called left-c.e. and right-c.e. with oracle $\emptyset^{(k-1)}$ respectively.

In Proposition 2.2, we saw that the hierarchy of real numbers can be reformulated using the hierarchy of subsets of $\mathbb{Q}$. It is also closely related to the hierarchy of closed subsets of $\mathbb{R}$.

Proposition 2.6. A real number x is Π_k if and only if the closed set $(-\infty, x]$ is Π_k in $\mathbb{R}$. Symmetrically, a real number x is Σ_k if and only if the closed set $[x, +\infty)$ is Π_k in $\mathbb{R}$.

Definition 2.7 (Computable Function). Let (X, d_X, S_X) and (Y, d_Y, S_Y) be two computable metric spaces. A function $f : X \rightarrow Y$ is **computable** if the sets $(f^{-1}(B_k^Y))_{k \in \mathbb{N}}$ are uniformly effective open sets, i.e. if there exists a c.e. set $E \subseteq \mathbb{N}^2$ such that, for every $k \in \mathbb{N}$, we have $f^{-1}(B_k^Y) = \bigcup_{(k,l) \in E} B_l^X$.

As the preimage of each rational ball B_k^Y is an open set, and these balls form a basis of the topology on Y , it follows that a computable function is continuous.

There are several equivalent definitions of computable functions, expressing the existence of an algorithm that, fed with arbitrary approximations of x , produces arbitrary approximations of $f(x)$. In particular, if $f : X \rightarrow Y$ is a computable function, then for any computable point $x \in X$, $f(x) \in Y$ is a computable point as well (and this second property is slightly weaker than the computability of f).

It turns out that, thus defined, computable functions interact with the effective versions of classical topological notions (open sets, compact sets...) in the same way continuous functions interact with those topological notions (we will see examples below). In Definition 2.7, we formulated computability as the effective version of continuity notably because it makes these relationships much more straightforward.

If $f : X \rightarrow Y$ is a computable function and $P \subseteq Y$ is Π_k , then its preimage $f^{-1}(P)$ is also Π_k . However, in general Π_k sets are not preserved by taking direct images, in particular because the image of a closed set by a continuous function is not always closed. However, we will see below that such a result holds under a computable compactness assumption about X (just like, for a continuous function, the image of a compact set is a compact set).

2.4 Computability of Compact and Closed Sets

Definition 2.8. Let (X, d, S) be a computable metric space. A compact subset $K \subseteq X$ is **computably compact** if its covering set $\{\langle k_0, \dots, k_n \rangle \in \mathbb{N}^* : K \subseteq B_{k_0} \cup \dots \cup B_{k_n}\}$ is c.e.

We say that the space X is computably compact if, as a subset of itself, it is a computably compact set. Notably, for a computably compact space X , given any positive rational radius r ,

we can compute in finite time a covering of X by rational balls of radius r (it is known to exist, so we only need to let the enumeration run long enough to encounter one). In a computably compact space, a set is computably compact if and only if it is Π_1 . The image of a computably compact set by a computable function is computably compact.

When the space is compact but not computably so, the complexity of the covering set of a Π_1 closed set (which is necessarily compact itself) may increase, but only from Σ_1^0 (c.e.) to Σ_2^0 .

Proposition 2.9. *If (X, d, S) is compact, then the covering set of any Π_1 closed set is Σ_2^0 .*

Proof. Let K a Π_1 closed set (i.e. $X \setminus K$ is effectively open). Let $(l_i)_{i \in \mathbb{N}}$ be a computable sequence of natural numbers such that $X \setminus K = \bigcup_{i \in \mathbb{N}} B_{l_i}$. For a rational ball $B_k = B(s, r)$ and a rational factor $a \in (0, 1)$, let $\bar{B}_k^a = \bar{B}(s, ar) = \{x \in X : d(s, x) \leq ar\}$. One has the following equivalences:

$$\begin{aligned}
K \subseteq B_{k_0} \cup \dots \cup B_{k_n} &\iff X \subseteq B_{k_0} \cup \dots \cup B_{k_n} \cup \bigcup_{i \in \mathbb{N}} B_{l_i} && (\text{as } X = K \sqcup X \setminus K) \\
&\iff \exists p, X \subseteq B_{k_0} \cup \dots \cup B_{k_n} \cup \bigcup_{i \leq p} B_{l_i} && (\text{by compactness}) \\
&\iff \exists p, a, X \subseteq \bar{B}_{k_0}^a \cup \dots \cup \bar{B}_{k_n}^a \cup \bigcup_{i \leq p} \bar{B}_{l_i}^a && (\text{by compactness}) \\
&\iff \exists p, a, \forall i, s_i \in \bar{B}_{k_0}^a \cup \dots \cup \bar{B}_{k_n}^a \cup \bigcup_{i \leq p} \bar{B}_{l_i}^a. && (\text{by density of } (s_i)_{i \in \mathbb{N}})
\end{aligned}$$

Observe now that, as the real numbers $d(s_i, s_j)$ are uniformly computable, the predicate $s_i \in \bar{B}_k^a$ is Π_1^0 (i.e. it can be replaced by some formula $\forall t, g(a, i, k, t)$ with g a computable algorithm). It follows that the predicate in the previous equivalence is itself Σ_2^0 . $\square$

Hierarchies of subsets of $\mathbb{N}$ and of real numbers are symmetric, in the sense that each class Π_k^0 or Π^k comes with its dual notion Σ_k^0 or Σ_k .

In practice, it turns out that the useful “dual” notion of Π_k for closed sets is not, as one would naturally expect, that of a Σ_k closed set (nor an analogous definition for open sets). Intuitively, a closed set is Π_1 if its complement can be enumerated as a union of rational balls. It does not make sense to enumerate the closed set itself, which is usually uncountable, but the next notion turns out to be a useful way of describing the set “from inside”.

Definition 2.10. *Let (X, d, S) be a computable metric space (X, d, S) . A closed set $A \subseteq X$ is **computably overt** if the set $\{k : A \cap B_k \neq \emptyset\}$ is c.e.*

Under the assumption that the metric is complete, computable overttness can be characterized in a particularly simple way [BP03, Theorem 3.8].

Proposition 2.11. *When the metric is complete, a closed set is computably overt if and only if it contains a uniformly computable dense sequence.*

The next result summarises the computable complexity of maximising a function over a compact set. These results appear in many places in the literature on computable analysis, but we include a sketch of proof for clarity.

Proposition 2.12. *Let $K \subseteq \mathbb{R}$ be compact:*

- If K is computably compact, then $\max K$ is Π_1 ,
- If K is computably overt, then $\max K$ is Σ_1 .

Let (X, d, S) a computable metric space, $f : X \rightarrow \mathbb{R}$ computable and $P \subseteq X$ compact.

- If P is Π_1 , then $\max f$ is Π_2 .
- If P is computably compact, then $\max f$ is Π_1 ,
- If P is computably overt, then $\max f$ is Σ_1 ,
- If P is both computably compact and overt, then $\max f$ is Δ_1 , i.e. computable.

Proof. One needs to show that the set $\{q \in \mathbb{Q} : \max K < q\}$ is c.e. to prove the first item. As K is compact (thus bounded), we have $\max K < q$ if and only if there exists a rational number $p < q$ such that $K \subseteq (p, q)$. This relation is c.e. when K is computably compact, so $\max K$ is Π_1 per Proposition 2.6.

If K is computably overt then it contains a dense computable sequence by Proposition 2.11, its supremum is Σ_1 by definition, and that supremum is precisely $\max K$.

Now, if P is Π_1 , then its covering set is Σ_2^0 by Proposition 2.9, and we must show that the covering set of $f(P)$ is also Σ_2^0 . Let $E \subseteq \mathbb{N}^2$ be a c.e. set such that $f^{-1}(B_k^{\mathbb{R}}) = \bigcup_{(k,l) \in E} B_l^X$. The sets $E_k = \{l : (k, l) \in E\}$ are uniformly c.e. By compactness of P , one has:

$$\begin{aligned} f(P) \subseteq B_{k_0}^{\mathbb{R}} \cup \dots \cup B_{k_n}^{\mathbb{R}} &\iff P \subseteq \bigcup_{l \in E_{k_0}} B_l^X \cup \dots \cup \bigcup_{l \in E_{k_n}} B_l^X \\ &\iff \exists l_0, \dots, l_p \in \bigcup_{0 \leq i \leq n} E_{k_i} \text{ such that } P \subseteq \bigcup_{0 \leq j \leq p} B_{l_j}^X, \end{aligned}$$

which is Σ_2^0 if the covering set of P is. Therefore, $\max f(P) < q$ if and only if $\exists p, f(P) \subseteq (p, q)$ which is Σ_2^0 , so $\max f(P)$ is Π_2 .

The remaining statements follow from the first two assertions, using the fact that taking the image of a set by a computable function preserves computable compactness and computable overtness. $\square$

2.5 Space of Probability Measures

When X is compact metrisable, the space $\mathcal{M}(X)$ of Borel probability measures over X endowed with the weak-* topology is metrisable, and a compatible metric is given by the Wasserstein metric

$$W_1(\mu, \nu) = \sup_{f \in \text{Lip}} \left| \int f d\mu - \int f d\nu \right|$$

where Lip is the set of 1-Lipschitz functions $f : X \rightarrow \mathbb{R}$ [AGS08, Section 7.1]. The separability of X implies the separability of $\mathcal{M}(X)$, where a dense sequence of measures is given by the set D of finite convex combinations with rational coefficients of Dirac measures supported on a countable dense subset of X . The compactness of X also implies the compactness of $\mathcal{M}(X)$ respectively. These implications also hold in a computable setting (see [HR09, Proposition 4.1.3] and [GHR11, Lemma 2.12]).

Proposition 2.13. *If (X, d, S) is a computable metric space, then $(\mathcal{M}(X), W_1, D)$ is also a computable metric space. If X is moreover computably compact, then so is $\mathcal{M}(X)$.*

The computable structure on $\mathcal{M}(X)$ makes integration of computable functions a computable operator [GHR11, Corollary 4.3.2].

Proposition 2.14. *If $\varphi : X \rightarrow \mathbb{R}$ is computable, then the function $\varphi^* : \mu \in \mathcal{M}(X) \mapsto \int \varphi d\mu \in \mathbb{R}$ is also computable.*

When X is computably compact, this result is uniform in φ , otherwise one needs to know an upper bound on $|\varphi|$ for each φ , and encode it in the corresponding algorithm.

If $T : X \rightarrow X$ is Borel measurable, then let $\mathcal{M}_T(X)$ be the set of T -invariant Borel probability measures over X .

Proposition 2.15. *Let (X, d, S) be a computable metric space and $T : X \rightarrow X$ be computable. The set $\mathcal{M}_T(X)$ is a Π_1 subset of $\mathcal{M}(X)$.*

The set $\mathcal{M}_T(X)$ is not computably overt in general. It was shown in [GHR11] that there exists a computable function from the circle to itself that has no computable invariant measure, so in that case $\mathcal{M}_T(X)$ does not contain any computable point.

3 General Theoretical Upper Bounds

Definition 3.1 (Computable Dynamical System). *A computable dynamical system (X, d, S, T) is, expectedly, a computable metric space (X, d, S) provided with a computable transformation $T : X \rightarrow X$.*

In the context of *ergodic optimisation*, our default assumption here is that X is compact (but *not necessarily* computably so) and that the potential $\varphi : X \rightarrow \mathbb{R}$ is computable. Consider the following hypotheses:

- ① X (and thus incidentally $\mathcal{M}(X)$ and $\mathcal{M}_T(X)$) is computably compact.
- ② $\mathcal{M}_T(X)$ is computably overt.

As explained in a general setting in the previous section, ① means that the compactness of X can be expressed (thus used) computationally. In particular, $\mathcal{M}_T(X)$ can be described “from the outside” by eliminating non-intersecting balls. This property usually follows from the fact that the space X has a “nice” and explicit definition.

Hypothesis ② is more complex to obtain (see Section 4 for a conversation on this topic in the context of ergodic optimisation) and implies that $\mathcal{M}_T(X)$ can itself be seen as a computable space, computably endowed with its own countable dense family.

Theorem 3.2. *Let (X, d, S, T) be a computable dynamical system, with X compact, and let $\varphi : X \rightarrow \mathbb{R}$ a computable potential. The following complexity upper bounds hold:*

Hypotheses	$\beta(\varphi)$ in $\mathbb{R}$	$\mathcal{M}_{\max}(\varphi)$ as a compact in $\mathcal{M}(X)$
General Case	Π_2	Π_3
① (X is comput. compact)	Π_1	Π_2
② ($\mathcal{M}_T(X)$ is comput. overt)	Σ_1	Π_1
① + ②	Δ_1	Π_1

Proof. As φ is computable, $\varphi^* : \mu \in \mathcal{M}(X) \mapsto \int \varphi d\mu$ is itself computable by Proposition 2.14.

The real number $\beta(\varphi)$ is thus the maximum of φ^* on the subset $\mathcal{M}_T(X)$ (which is Π_1 from Proposition 2.15), so it is Π_2 (from Proposition 2.12). The other bounds on $\beta(\varphi)$ follow by likewise applying Proposition 2.12 when $\mathcal{M}_T(X)$ is computably compact and/or overt.

Now, as $\mathcal{M}_T(X)$ itself is a Π_1 set (Proposition 2.15), it suffices to obtain a Π_k bound on $\{\mu \in \mathcal{M}(X) : \int \varphi d\mu \geq \beta(\varphi)\} = \varphi^{*-1}([\beta(\varphi), \infty))$ for it to apply to $\mathcal{M}_{\max}(\varphi)$ too. As φ^* is computable, the preimage of a Π_k closed set is itself a Π_k closed set of $\mathcal{M}(X)$. To obtain a bound, we thus simply need to find k such that $[\beta(\varphi), \infty)$ is a Π_k subset of $\mathbb{R}$, *i.e.* $\beta(\varphi)$ is a Σ_k real number (see Proposition 2.6). Noticing $\Pi_k \subset \Sigma_{k+1}$ (we simply add a leftmost useless inf), the result follows from the first part of the theorem. $\square$

4 Brief Survey of the Complexity of Dynamical Systems

To ensure our study and hierarchy in Theorem 3.2 is, to a degree, rooted in reality, we will in this section list the dynamical systems most representative of (the early days of) ergodic optimisation, and discuss whether they satisfy ① and/or ②, whether they are computably compact and/or overt.

This section is purposefully an informal conversation, it includes no result that we consider new, though some facts are hard-to-pinpoint folklore.

4.1 General Context

As already explained, ① directly follows from the definition of the space X and thus is easy to satisfy on “naturally occurring” spaces (and it will be satisfied on all the examples we discuss), whereas ② is a convoluted function of how T acts on X , not always satisfied and harder to establish – positively or negatively. Still, those two notions are theoretically unrelated and complementary, and there are dynamical systems for which only ① or ② (or none) holds.

We acknowledge the focus of ergodic optimisation in the past twenty years has shifted from the compact framework we use to non-compact settings (notably by having either an unbounded real space such as $\mathbb{R}^d$, or a shift space $\mathcal{A}^G$ on a non-compact alphabet such as $\mathcal{A} = \mathbb{N}$), which fall outside of the established scope of the paper, making ② notably out of reach.

Discontinuous transformations (such as the Gauss map $T : x \mapsto \frac{1}{x} - \lfloor \frac{1}{x} \rfloor$ on the unit interval) are also much more intrinsically out of scope, by virtue of being non-computable.

Another prevalent case that escapes the scope of this paper is that of continuous flows where, instead of a computable transformation $T : X \rightarrow X$, we have a family $T = (T_t)_{t \in \mathbb{R}^+}$ (such that $T_u \circ T_v = T_{u+v}$ for any $u, v \geq 0$, with $(t, x) \mapsto T_t(x)$ computable). This plays no role on whether ① holds, but makes the study of ② even harder – however, assuming ② holds for the set of measures invariant for *the whole family* T , the complexity upper bounds in the previous section should follow likewise.

4.2 One Dynamic in a Three-Models-Shaped Trenchcoat

Consider $P(x) = 4x(1 - x)$ the Ulam-von Neumann map on $X = [0, 1]$ (with the usual metric on $\mathbb{R}$), the expanding map $T(\theta) = 2\theta$ on $Y = \mathbb{R}/\mathbb{Z}$ (with the corresponding induced metric) and the full shift σ on $Z = \{0, 1\}^{\mathbb{N}}$ (with $d(x, y) = 2^{-\inf\{i : x_i \neq y_i\}}$). All these models were mentioned

in Jenkinson’s 2006 lecture notes [Jen06] (also available [on their webpage](#)), which makes them staple examples for ergodic optimisation.

Due to its explicit symbolic structure, the full shift (Z, σ, d) is a computationally well-behaved space, with the dense family $(w0^\infty)_{w \in \{0,1\}^*}$. The rational balls are given by the clopen cylinders $([w])_{w \in \{0,1\}^*}$, so that the covering set is not only c.e. but *decidable*, ① holds trivially. For periodic configurations $\omega = w^\infty$ (with $w \in \{0,1\}^*$ a finite word of length $p = |w|$), the corresponding “periodic measures” $\mu_\omega = \frac{1}{p} \sum_{i=0}^{p-1} \delta_{\sigma^i(\omega)} \in \mathcal{M}_\sigma(Z)$ represent a uniformly computable dense sequence, so that ② holds.

Consider $\Gamma : Z \rightarrow Y$ defined as $\omega \mapsto \sum_{i \in \mathbb{N}} \frac{\omega_i}{2^i}$. The map is not injective, as collisions between exactly two elements occur for any words of the form $u01^\infty$ and $u10^\infty$ (with a prefix $u \in \{0,1\}^*$), as well as 1^∞ and 0^∞ . This set of words U can only support exactly two σ -invariant measures, on the fixed points 0^∞ and 1^∞ . Likewise, looking at the image set $\Gamma(U)$, we obtain specifically dyadic fractions, which all get sent to the fixed point 0 after finitely many iterations. Outside of those, $\Gamma : Z \setminus U \rightarrow Y \setminus \Gamma(U)$ is a bimeasurable bijection such that $\Gamma \circ \sigma = T \circ \Gamma$, establishing notably a one-to-one correspondence between the invariant measures, and computably matches the (dense) periodic sequences (*i.e.* a computation basis of Z) and (dense) non-dyadic rational numbers (*i.e.* a computation basis of Y).

Likewise, with $\Psi : Y \rightarrow X$ defined as $\theta \mapsto \sin^2(\pi\theta)$, we have a non-injective computable map with two branches (it computably bijects $[0, \frac{1}{2}]$ with $[0, 1]$ going upwards, and then $[\frac{1}{2}, 1]$ with $[0, 1]$ going downwards, the only one-to-one points being $\Psi(0) = 0$ and $\Psi(\frac{1}{2}) = 1$). Notably, $\Psi(x) = \Psi(-x) = \Psi(1-x)$. Now, $\Psi \circ T = P \circ \Psi$, so once again, they both encode in superficially different ways the same core dynamical behaviour, and the dense family of periodic orbits in Y thus gives a dense family of periodic orbits in X .

The point of all this is that the full shift Z is much easier to study computably to obtain ②, which in turn implies that ② also holds for Y and X in a very straightforward way, much simpler than trying to find the uniformly computable and dense periodic orbits in X directly (*i.e.* by solving the polynomial equations $P \circ \dots \circ P(x) = x$) and identifying the structure of $\mathcal{M}_P(X)$ itself.

This correspondence also allows us to exhibit “weird” invariant measures for (X, P) or (Y, T) corresponding to known atypical ones on the full shift, such as the (unique) invariant measure on the strictly ergodic Thue-Morse subshift.

4.3 Other Models

Let us mention some more models, taken from Jenkinson’s more recent survey [Jen19].

Generalising the aforementioned expansion $T_0 : x \mapsto 2x$ of Y (on which ① holds), there is a whole family of Manneville-Pomeau expanding circle maps of the form $T_\alpha : x \mapsto x + x^{1+\alpha}$ (with $\alpha \geq 0$). In this case, there is still a (topologically semi-conjugate) full shift hidden inside, as is more broadly the case for *any* expanding map on Y of any degree $d \geq 2$ [KH03, Section 7.4]. Identifying this full shift means identifying the *branches* of the expanding map (whose endpoints are the preimages of 0). For instance, as T_α is of degree 2 here, it would mean computing the unique $0 < \gamma < 1$ such that $\gamma + \gamma^{1+\alpha} = 1$, by dichotomic search for example. We can then repeat this process by computing the two pre-images of γ , on and on to obtain the 2^k branches of $T_\alpha^k := T_\alpha \circ \dots \circ T_\alpha$ k times. By excluding this countable set C of uniformly computable points (analogous to dyadic fractions in the case of T_0 , whose orbits eventually sink in the fixed point 0), we can

thus simply map the remaining elements $x \in Y \setminus C \mapsto (\mathbb{1}_{(\gamma,1)}(T_a^k(x))) \in \{0,1\}^{\mathbb{N}}$, establishing by design a computable correspondence between invariant measures of the binary full shift (Z, σ) and those of any Manneville-Pomeau circle map (Y, T_a) , provided a is itself computable (and more broadly between the full shift $[[1, d]]^{\mathbb{N}}$ and any computable expanding map of degree d), which guarantees ② holds.

Another interesting example is the quintessential model of spin glasses, the XY model, in which we don't just look at *one* spin in Y , but a whole glass in $Y^{\mathbb{Z}}$, with the shift σ as a natural transformation. In this case, periodic configurations with “rational” spins in $\mathbb{Q}/\mathbb{Z}$ (or rather their projection on Y) provide both a computably enumerable dense family of $Y^{\mathbb{Z}}$ (hence define a notion of computability) and an induced dense family of invariant measures (so that ② holds). By using the same computable compactness as for Y (with finite coverings by increasingly small rational intervals B_k) but on increasingly large windows $[-n, n] \subset \mathbb{Z}$ of spins, one can easily reach the conclusion that ① still holds (by checking, for each potential covering $(B_{\langle k^i \rangle})_i$ of $Y^{\mathbb{Z}}$ by “rational rectangles” $B_{\langle k \rangle} := Y^{\infty} \times B_{k_{-n}} \times \cdots \times B_{k_n} \times Y^{\infty}$, if we can find a fine enough covering of Y , extended as a covering of $Y^{[-n, n]}$ by small “rational cubes”, that refines $(B_{\langle k^i \rangle})_i$).

All the examples discussed so far are deeply related to a full shift space, but this is not necessarily the case in general.

4.4 Generalising the Ulam-von Neumann Map

Generalising the Ulam-von Neumann map gives us the family of logistic maps $f_{\lambda} : x \mapsto \lambda x(1-x)$ for $\lambda \in [0, 4]$.

These maps are part of a broader family of continuous interval transformations $T : I \rightarrow I$, for which highly chaotic behaviours can occur, notably regarding which periodic orbits can exist (*Period Three Implies Chaos* [LY75], and more broadly the existence of p -periodic orbits implies the existence of q -periodic orbits for any $p \succcurlyeq q$ in *Sharkovskii's order* [Sha64; Sha95]).

In the case of logistic maps f_{λ} , from $\lambda \leq 1$ (where 0 is the unique attractor) to $\lambda > 1 + \sqrt{8}$ (where f_{λ} admits a 3-periodic orbit, thus of any other period), we cross many thresholds that change the maximal period (in Sharkovskii's order). Notably, starting from $\lambda \leq 1$, we encounter a series of period-doubling thresholds (two-periodic orbits appear beyond $\lambda = 3$, 4-periodic orbits appear beyond $\lambda = 1 + \sqrt{6}$, accumulating (at a geometric rate known as the *Feigenbaum constant*) to a computable limit [Hoy07, Theorem 6] for which no proper algebraic expression is known.

Some of these thresholds correspond to the apparition of isolated periodic points, thus represent discontinuities in the map $\lambda \mapsto \mathcal{M}_{\max}(f_{\lambda})$. For the usual reasons of undecidability of equality over computable real numbers, unless we are given supplementary information on an *a priori* arbitrary computable number λ , it could lie on such a problematic threshold, in which case we will be unable to properly satisfy ② (even if the threshold λ_t itself is a computable real number and we mathematically know that $\mathcal{M}_{\max}(f_{\lambda_t})$ is computably overt).

A likewise potential obstruction to ② can be found in the existence of computable parameters λ for which f_{λ} has a non-computable SRB measure [RY20; RY24].

4.5 Symbolic Shift Spaces

For simplicity's sake, we will say a shift space X is a (closed) subset of $\mathcal{A}^{\mathbb{N}}$ (with a finite alphabet $\mathcal{A}$) invariant under the shift action σ . Symbolic shift spaces are “everywhere”, beyond the works

cited by Jenkinson’s (extensive) surveys, for many reasons we will not detail. They will become our focal point in the following final section of this article.

As seen above a full shift can, with its very straightforward dynamics, deobfuscate the complex behaviour behind analytic systems such as the Ulam-von Neumann map $x \mapsto 4x(1 - x)$ on the unit interval.

Given a set of finite words $\mathcal{F} \subset \mathcal{A}^*$ called *forbidden patterns*, one can define the induced subshift $X_{\mathcal{F}}$ made of configurations avoiding $\mathcal{F}$. For any subshift X , denote $L(X) \subset \mathcal{A}^*$ its *language* (each word in $L(X)$ is the finite prefix of an element of X), then the countable set $\mathcal{F} = \mathcal{A}^* \setminus L(X)$ is such that $X = X_{\mathcal{F}}$.

This apparent structural simplicity of the full shift hides the very complex and rich behaviours that can occur on subshifts depending on the setting, whether we are looking at a one-sided or two-sided shift (in $\mathcal{A}^{\mathbb{Z}}$), whether it is a *Subshift of Finite Type* (an SFT, such that $X = X_{\mathcal{F}}$ for a finite set $\mathcal{F}$), sofic (with $\mathcal{F}$ a regular language) or effective (with $\mathcal{F}$ a c.e. subset of $\mathcal{A}^*$).

Another point of interest, which we will briefly discuss, is the case of *higher-dimensional* subshifts, *i.e.* closed subsets of $\mathcal{A}^{\mathbb{Z}^d}$ invariant under translations along every direction. Of interest here is the fact that for finite sets of forbidden patterns $\mathcal{F}$, *unlike* the one-dimensional case where there is an algorithmic procedure to know whether $X_{\mathcal{F}}$ is non-empty (and if so, then we can computably enumerate increasingly long periodic orbits which give a dense family of invariant measures, such that ② holds), the higher-dimensional domino problem is *undecidable* and there are explicit non-empty SFTs (such as the Robinson or the Kari-Čulík tiling) for which no periodic orbit exists or, even worse, where $\mathcal{F}$ is explicit but not a single point of $X_{\mathcal{F}}$ is a computable point of $\mathcal{A}^{\mathbb{Z}^2}$ [Han74; Mye74] (so that even endowing it with a computable space structure would require at best careful work, before starting to discuss ① and ②). Consequently, the periodic orbits can no longer form an easily enumerable dense family to satisfy ②, as was the case in all the previous examples – this argument might appear a bit fallacious given the Robinson tiling is uniquely ergodic (thus trivially satisfies ②) and its invariant measure is very well known and understood, but the point we are making here is that in the general two-dimensional case, it is *impossible* to algorithmically obtain this information on $\mathcal{M}_T(X_{\mathcal{F}})$ being given $\mathcal{F}$, and thus this provides a very natural class of dynamical systems *simple to describe but hard to understand* for which ② is unknown in general.

Back to the one-sided one-dimensional case, in the following section, using graph arguments with a flavour similar to that of those used to decide whether $X_{\mathcal{F}}$ is non-empty and if so to find its periodic orbits, we will give algorithmic arguments to compute both $\beta(\varphi)$ and $\mathcal{M}_{\max}(\varphi)$ for finite-range rational potentials $\varphi : \mathcal{A}^r \rightarrow \mathbb{Q}$.

5 Finite-Time Polynomial Complexity for Markov Shifts

In Theorem 3.2, we identified a setting in which $\beta(\varphi)$ is computable.

By assuming the space X is computably compact (①), we can enumerate arbitrarily accurate finite coverings of $\mathcal{M}_T(X)$ by rational balls, and each such covering gives us an accordingly close upper bound on $\beta(\varphi)$. By assuming $\mathcal{M}_T(X)$ is computably overt (②), we obtain a uniformly computable, countable, dense subset $S \subset \mathcal{M}_T(X)$ (*e.g.* periodic orbits for a full shift), and getting arbitrarily good approximations of $\varphi^*(\mu)$ for each $\mu \in S$ gives an arbitrarily close lower bound on $\beta(\varphi)$. Putting those two together, assuming ①+②, we obtain an algorithm that can

approximate $\beta(\varphi)$ with arbitrary precision.

The main drawback of this all-purposes approach is that we have no specific information on the structure of (X, T) as a dynamical system, everything is abstracted away and obfuscated behind our computable objects. Consequently, we can provide no information on the *computational complexity* of the method.

The goal of this final section is precisely to address this blind spot by considering a very explicit model both computationally simple to manipulate and mathematically well-understood, that of one-dimensional SFTs, sometimes also called *Markov shifts* in the mathematical literature.

5.1 Controlling the Maximising Measures for Well-Behaved Potentials

Let (X, T) be a compact dynamical system, with $\varphi : X \rightarrow \mathbb{R}$ a continuous potential. Intuitively, what is $\mathcal{M}_{\max}(\varphi)$ expected to look like?

Assuming φ is constant on every T -orbit (which is a way too strong assumption), it is quite easy to see that $\beta = \max \varphi$ and that $\mathcal{M}_{\max}(\varphi) = \mathcal{M}_T(\varphi^{-1}(\beta))$. More reasonably, as φ gives a weight to every configuration, it informally induces an average weight for every orbit. By defining Y as the union of all the φ -maximal orbits, one can expect to have $\mathcal{M}_{\max}(\varphi) = \mathcal{M}_T(Y)$, *i.e.* the maximising measures are those supported by maximising orbits.

While this statement is not provably true in a general setting, it actually holds for some well-behaved dynamical systems (*e.g.* weakly expanding like the one-sided shift $\mathcal{A}^{\mathbb{N}}$, or hyperbolic with weak local product like the two-sided shift $\mathcal{A}^{\mathbb{Z}}$ [Bou01]), provided φ is regular-enough (it must satisfy the *Walters condition* so that some version of *Mañé's lemma* holds).

The Walters condition on a continuous potential $\varphi : X \rightarrow \mathbb{R}$, as defined in Bousch's article [Bou01], is a sort of uniform continuity property for φ along trajectories of arbitrary lengths. More precisely, with (X, d, T) a metric compact system, if the T -orbits of $x, y \in X$ stay ε -close for some (arbitrary) amount $n \in \mathbb{N}$ of T -transitions, then

$$\left| \sum_{i=0}^n \varphi \circ T^i(x) - \varphi \circ T^i(y) \right| \leq h(\varepsilon),$$

with $h : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ a non-decreasing function independent of n , such that $\lim_0 h = 0$, called a *Walters module* for φ .

Remark 5.1. *The notion of Walters module is independent from the computability of a potential φ . Trivially, non-computable finite-range (thus Walters) potentials exist. Less obviously but still easily, we can exhibit hand-crafted non-Walters computable potentials for which $\beta(\varphi)$ and $\mathcal{M}_{\sigma}(\varphi)$ are explicit.*

Notably, in the case of a one-sided shift space $\mathcal{A}^{\mathbb{N}}$ (though every argument translates for two-sided shifts), with $d(x, y) = |\mathcal{A}|^{-\min\{i \in \mathbb{N} : x_i \neq y_i\}}$, the orbits of x and y being ε -close for n steps actually means they were $\frac{\varepsilon}{|\mathcal{A}|^i}$ -close at step $n - i$. Thus, when φ has summable variations (*i.e.* $F : \varepsilon \mapsto \sup_{d(x, y) \leq \varepsilon} |\varphi(x) - \varphi(y)|$ is integrable on a neighbourhood of 0), $\eta \mapsto \int_0^\eta F$ provides a Walters module. This argument holds in particular for finite-range potentials, which are *locally constant*.

This class of interactions has been widely studied, and is heavily related to SFTs. Notably, using cohomology arguments, Bousch proved that if φ is a finite-range potential, then there is

an SFT X such that $\mathcal{M}_{\max}(\varphi) = \mathcal{M}_{\sigma}(X)$. What he did *not* explain is how one can construct, or *compute*, a set $\mathcal{F}$ of forbidden patterns such that $X = X_{\mathcal{F}}$. This will be our focal point from now on.

Remark 5.2. *There are fundamental obstructions to an exact computation of $\mathcal{F}$ in finite time, due to the undecidability of equality over computable real numbers. If we consider a general computable $\varphi : \mathcal{A}^r \rightarrow \mathbb{R}$, we cannot have an algorithm providing $\mathcal{F}$ in finite time.*

Consider the following counter-example, on the binary alphabet $\mathcal{A} = \{0, 1\}$. Let $\varphi : \mathcal{A}^2 \rightarrow \mathbb{R}$ such that $\varphi(01) = \varphi(10) = -1, \varphi(00) = 0$ and $\varphi(11) = \lambda$ is some computable real number. Denote X the SFT that supports maximising measures. Switching from 0 to 1 in a configuration in $\mathcal{A}^{\mathbb{N}}$ always results in a clear negative penalty compared to $\beta = \max(0, \lambda)$, corresponding to the fixed-point configurations 0^{∞} and/or 1^{∞} . In this case, either $\lambda > 0$ (so $X = \{1^{\infty}\}$), $\lambda < 0$ (so $X = \{0^{\infty}\}$) or $\lambda = 0$ (so $X = \{0^{\infty}, 1^{\infty}\}$). While $\lambda \neq 0$ can be observed in finite time, the case $\lambda = 0$ can only be de facto concluded after infinitely many computations, and even providing a subset of X , or providing some measure $\mu \in \mathcal{M}_{\max}(\varphi)$, cannot be done in finite time for the same reasons.

Using this counterexample, we can see how impossible it is to exhibit even one measure in $\mathcal{M}_{\max}(\varphi)$. After finitely many computation steps, the most we can know mathematically about λ is an open rational interval I such that $\lambda \in I$. For a given such interval I , assuming $0 \in I$ too, any measure supported by $\{0^{\infty}, 1^{\infty}\}$ could be in $\mathcal{M}_{\max}(\varphi)$ (when $\lambda = 0$), but I also matches mutually exclusive cases with $\mathcal{M}_{\max}(\varphi)$ being $\{\delta_{0^{\infty}}\}$ if $\lambda < 0$, or $\{\delta_{1^{\infty}}\}$ if $\lambda > 0$. Any measure μ could be a wrong guess, and the interval I itself could be arbitrarily small so there is no hope for some sort of approximation $I \mapsto \mu_I$ such that $d(\mu_I, \mathcal{M}_{\max}(\varphi)) \rightarrow 0$ as $\text{diam}(I) \rightarrow 0$. This computability obstruction is intrinsic, it cannot be fixed by algorithmic means unless we are provided with an “oracle” telling us whether any two computable numbers are equal.

The best we could do in such a setting, adapting the ideas and algorithms that will be introduced below, is to replace each computable number used in φ by an approximate rational interval of length $\frac{1}{2^n}$, and eliminate the configurations that cannot be in X , that cannot maximise $\varphi^(\mu)$, rather than trying to find one that is. This would give us an algorithm $(\varphi, n) \mapsto \mathcal{F}_n \subset \mathcal{A}^r$ such that $\mathcal{M}_{\max}(\varphi) \subset \mathcal{M}_{\sigma}(X_{\mathcal{F}_n})$. In this scenario, the set of forbidden patterns $\mathcal{F}_n$ is increasing, stationary as $\mathcal{A}^r$ is finite, and converges to the “true” $\mathcal{F}$ such that $X_{\mathcal{F}}$ is the support of maximising measures. This fact was already a consequence of Theorem 3.2, as being computed “from the outside” by elimination is exactly what Π_1 means, and we are in a setting where ①+② holds.*

Hence, to provide an algorithm computing $\varphi \mapsto \mathcal{F}$ (so that $\mathcal{M}_{\max}(\varphi) = \mathcal{M}_{\sigma}(X_{\mathcal{F}})$), we will need some further assumptions on the values taken by φ , and clarify what we mean here by algorithm and complexity.

5.2 From Cohomology to Graph Algorithms

In the following, we consider $(\mathbb{F}, +, \times)$ an abstract ordered subfield of $\mathbb{R}$. The field $\mathbb{F}$ could simply be $\mathbb{R}$, or the subfield of all computable real numbers, but as discussed in the previous remark this would evade our “practical computations” approach. Thus, in the following algorithmic considerations, we assume that we can “compute” the sum $(x, y) \mapsto x + y$, the rational product $(q, x) \mapsto q \times x$ (with $q \in \mathbb{Q} \subset \mathbb{F}$) and test equality $(x, y) \mapsto (x = y)$ and inequality $(x, y) \mapsto (x < y)$

for our elementary operations. This holds for $\mathbb{Q}$ (by encoding a rational as pair of integers), but also (in a very unpractical way) for the field of all real algebraic numbers.

Remark 5.3. *The expression “elementary operations” is a bit deceitful here, as one would expect our individual variables in $\mathbb{F}$ to take $O(1)$ space, and our operations to take $O(1)$ time, but this obviously cannot hold for an infinite set such as $\mathbb{Q}$.*

It must be understood as already-implemented black boxes, which we will use as basic blocks of the algorithm, and which we will count to describe the time/space complexity of said algorithm. In practical implementations, we must choose between using $O(1)$ space/time objects (such as 64bit floating-point numbers) with rounding errors (breaking $\varphi \mapsto \mathcal{F}$ around fringe cases), or a dynamically-sized structure (such as Python long integers and fractions) with perfect precision but operations whose complexity depends on the size of the input.

We will consider a finite-range potential $\varphi : \mathcal{A}^{\mathbb{N}} \rightarrow \mathbb{F}$. As φ is locally constant, there is a cylinder size r such that, on any cylinder $[w]$ with $w \in \mathcal{A}^r$, $\varphi([w])$ is a singleton. We can thus equivalently see this potential as a function $\varphi : \mathcal{A}^r \rightarrow \mathbb{F}$, and this is how it will be “encoded” in a computer, by hardcoding its range r in its definition.

A useful tool, both for the next proposition and for what follows afterwards, will be the *De Bruijn* directed graph G_n on words of length $n > 0$, with transitions between words overlapping over $n - 1$ letters, i.e. $G_n := (\mathcal{A}^n, \{aw \rightarrow wb : a, b \in \mathcal{A}, w \in \mathcal{A}^{n-1}\})$.

Proposition 5.4. *Let $(\mathcal{A}^{\mathbb{N}}, \sigma)$ be the one-sided full shift on $\mathcal{A}$. Let $\varphi : \mathcal{A}^r \rightarrow \mathbb{F}$ be a potential of finite range $r \in \mathbb{N}$ (so that $\varphi(x) := \varphi(x_0 \dots x_{r-1})$ for an actual infinite word $x \in \mathcal{A}^{\mathbb{N}}$), and X be the (smallest) SFT such that $\mathcal{M}_{\max}(\varphi) = \mathcal{M}_{\sigma}(X)$ (i.e. it is, formally, the closure of the union of the supports of the measures in $\mathcal{M}_{\max}(\varphi)$, though there must necessarily exist a measure in $\mathcal{M}_{\max}(\varphi)$ fully supported on X).*

Then a finite word satisfies $w \in L(X)$ (i.e. it is the prefix of some element $x \in X$) if and only if there exists a periodic word $\omega \in [w]$ for which $\overline{\varphi}(\omega)$ – the (well-defined) average value of φ over ω – satisfies $\overline{\varphi}(\omega) = \beta(\varphi) \in \mathbb{F}$.

Proof. The $(\Leftarrow)$ implication is direct: if we have $w \in \mathcal{A}^*$ and $\omega \in [w] \subset \mathcal{A}^{\mathbb{N}}$ a p -periodic word satisfying $\overline{\varphi}(\omega) = \beta$, then by averaging over the finite σ -orbit of ω , we obtain the invariant measure $\mu_{\omega} = \frac{1}{p} \sum_{i=0}^{p-1} \delta_{\sigma^i(\omega)} \in \mathcal{M}_{\sigma}(\mathcal{A}^{\mathbb{N}})$ such that $\int \varphi d\mu_{\omega} = \overline{\varphi}(\omega) = \beta$, i.e. $\mu_{\omega} \in \mathcal{M}_{\max}(\varphi)$. As $\mu_{\omega}([w]) > 0$, it follows that w is in the language of the support of μ_{ω} hence in X .

For the $(\Rightarrow)$ implication, we will need to introduce an auxiliary function *cohomologous* to φ . Following the footsteps of Thierry Bousch [Bou01, Théorème 1, Proposition 10], we know that φ is *cohomologous* to a likewise finite-range function $h : \mathcal{A}^r \rightarrow \mathbb{R}$ such that $h \leq \beta$.

Given the aforementioned article is written in French, let us give some further clarifications that might otherwise remain lost in translation. Here, *cohomologous* must be understood as the existence of a continuous map $u \in C^0(\mathcal{A}^{\mathbb{N}}, \mathbb{R})$ such that $h = \varphi + u - u \circ \sigma$. In this case, he obtains a u that satisfies the functional equation $u(x) = \max_{a \in \mathcal{A}} (\varphi - \beta + u)(ax)$. Denoting $\tilde{\varphi} = \varphi - \beta$, we obtain $u(x) = \max_{a_k, \dots, a_1 \in \mathcal{A}} \sum_{i=1}^k \tilde{\varphi}(a_i \dots a_1 x) + u(a_k \dots a_1 x)$ by iterating. As u is uniformly continuous on the compact $\mathcal{A}^{\mathbb{N}}$, for any $\varepsilon > 0$, we have a rank k such that $\sup_{w \in \mathcal{A}^k, x, y \in \mathcal{A}^{\mathbb{N}}} |u(wx) - u(wy)| \leq \varepsilon$. Furthermore, if x and y share the same first $r - 1$ symbols, then the term $\sum_{i=1}^k \tilde{\varphi}(\dots)$ is the same in the max for $u(x)$ and $u(y)$, in which case we can easily deduce that $u(x) = u(y)$ as $k \rightarrow \infty$. In other words, we can represent it as $u : \mathcal{A}^{r-1} \rightarrow \mathbb{R}$.

As mentioned by Bousch, if we denote $K = h^{-1}(\beta) \subset \mathcal{A}^{\mathbb{N}}$ (with $h : \mathcal{A}^{\mathbb{N}} \rightarrow \mathbb{R}$), then as $h \leq \beta$, for any (not necessarily invariant) probability measure μ , we have $\int h d\mu \leq \beta$, with an equality if and only if $h = \beta$ μ -almost-surely, that is if $\text{supp}(\mu) \subset K$. For invariant measures, we can furthermore say that $\mu \in \mathcal{M}_{\max}(h)$ if and only if $\text{supp}(\mu) \subset Y := \bigcap_{i \in \mathbb{N}} \sigma^{-i}(K)$. Now, a vertex w_0 of the De Bruijn graph G_r satisfies $w_0 \in L(Y)$ if and only if there exists an infinite path $(w_i)_{i \in \mathbb{N}}$ such that for any $i \in \mathbb{N}$, $h(w_i) = \beta$.

In particular, if we consider $\mathcal{F} = \mathcal{A}^r \setminus h^{-1}(\beta)$ (with $h : \mathcal{A}^r \rightarrow \mathbb{R}$), we obtain a set of forbidden patterns such that $Y = X_{\mathcal{F}}$ is precisely the SFT induced by $\mathcal{F}$ (and then by removing vertices/words that fall outside of a strongly connected component we would obtain X itself). From thereon, it would be expected that we compute h directly, and then define/describe X by restricting G_r to strongly connected components of h -maximising vertices. However, the main issue with this approach is that h itself is *known to exist* by compactness arguments, but the proof of its existence doesn't provide an algorithmic procedure to compute it exactly, not even as a limit with known speed of convergence. The end of the proof, and this whole proposition, aim at allowing us to work around this obstruction with a practical argument.

All this context being laid out, consider $w \in L(X)$, so that by ergodic decomposition, we have some *ergodic* measure $\mu \in \mathcal{M}_{\max}(\varphi)$ for which $\mu([w]) > 0$. By ergodicity, we have a configuration $\omega \in X \subset Y$ for which $\mu = \mu_{\omega}$ (i.e. it is the well-defined Cesàro sum of the sequence $(\delta_{\sigma^i(\omega)})_{i \in \mathbb{N}}$). In particular, $\mu([w])$ is the (positive) frequency at which $\sigma^i(\omega) \in [w]$. Equivalently, $\omega \in Y$ can be seen as an infinite path (w_i) on the h -maximising vertices of G_r . By finiteness of $\mathcal{A}^r$, we can fix two steps s, t such that $s + |w| < t$, $\sigma^s(\omega), \sigma^t(\omega) \in [w]$ and $w_s = w_t$. We then define the $(t - s)$ -periodic configuration $\omega' = (w_s, w_{s+1}, \dots, w_{t-1}, w_t = w_s, \dots)$. Now, w has a positive frequency in ω' (at least $\frac{1}{t-s}$), so $\mu_{\omega'}([w]) > 0$ still, and w is a prefix of ω' , i.e. $\omega' \in [w]$.

The final nail in the coffin is that, by construction:

- ω' is periodic, so that $\bar{h}(\omega') = \int h d\mu_{\omega'}$ and $\bar{\varphi}(\omega') = \int \varphi d\mu_{\omega'}$ are both well-defined,
- ω' only visits h -maximising vertices of G_r , so $\bar{h}(\omega') = \beta$,
- φ and h are cohomologous, so $\int \varphi d\mu_{\omega'} = \int h d\mu_{\omega'}$,

so we know that $\bar{\varphi}(\omega') = \beta$, which concludes the proof. $\square$

Corollary 5.5. *Let us denote $n = |\mathcal{A}^r| = |\mathcal{A}|^r$. For $w \in \mathcal{A}^r$, define $F(w)$ as the maximum of $\bar{\varphi}(\omega)$ among all the cycles (w_i) on G_r of period $p \leq n$ with $w_0 = w$. Then we have:*

- $F(w) \leq \beta$,
- $F(w) = \beta$ if and only if $w \in L(X)$,
- $X = X_{\mathcal{F}}$ with $\mathcal{F} = \mathcal{A}^r \setminus F^{-1}(\beta)$.

Proof. The first item follows from the fact that any cycle ω satisfies $\bar{\varphi}(\omega) = \int \varphi d\mu_{\omega} \leq \beta$.

For the second item, clearly $F(w) = \beta$ gives the existence of a cycle ω that guarantees $w \in L(X)$ in the previous proposition. Conversely, $w \in L(X)$ guarantees the existence of a h -maximising p -periodic cycle $\omega = (w_i)$, but it might be *too long*. However, if $p > n$, then ω contains a repetition, some ranks $0 \leq i < j < p$ such that $w_i = w_j$ and so, by skipping all that happens in-between, $(w_0, \dots, w_{i-1}, w_i = w_j, w_{j+1}, \dots, w_{p-1})$ gives a shorter h -maximising cycle such that $w_0 = w$ (notably, this property breaks for a general $w \in L(X)$: if it is too long to be

uniquely determined by w_0 , by shortening the cycle, we might break the $\omega \in [w]$ property). By induction, we can always obtain a short-enough period $p \leq n$, so that $F(w) = \overline{\varphi}(\omega) = \beta$.

For the third item, we simply need to remark that $F^{-1}(\beta) \subset \mathcal{A}^r$ is precisely the set $L(X) \cap \mathcal{A}^r$, made of all the vertices of G_r that belong to some strongly connected h -maximising component, so $X = X_{\mathcal{F}}$ directly follows. $\square$

Consequently, X can be characterised directly through an algorithmic study of φ itself, without involving any knowledge about the cohomologous function $h = \varphi + u - u \circ \sigma$.

Remark 5.6. *The previous proposition and corollary both consider a full shift $(\mathcal{A}^{\mathbb{N}}, \sigma)$ as the underlying dynamical system. One might wonder what happens if we consider instead an SFT $X_{\mathcal{G}}$ (defined through $\mathcal{A}$ and $\mathcal{G} \subset \mathcal{A}^*$) to maximise the potential among measures in $\mathcal{M}_{\sigma}(X_{\mathcal{G}})$.*

Without loss of generality, by adding “padding” symbols to the elements of $\mathcal{G}$ and/or φ , we might assume that $\mathcal{G} \subset \mathcal{A}^r$ and $\varphi : \mathcal{A}^r \rightarrow \mathbb{F}$.

Then the previous results apply likewise by considering instead periodic words defined on the restriction of the De Bruijn graph G_r induced on the subset of vertices $\mathcal{A}^r \setminus \mathcal{G}$.

5.3 Explicit Algorithms for the Maximum Mean Weight Cycles

From the previous analysis, it follows that in the case of a finite-range potential $\varphi : \mathcal{A}^r \rightarrow \mathbb{F}$, computing β and $\mathcal{F}$ amounts to computing the average weights of cycles of a directed graph with the vertices $\mathcal{A}^r$.

In the case of a strongly connected directed graph $G = (V, E)$ with weights on the *edges*, Karp’s algorithm [Kar78] is known to compute the maximal mean weight of a cycle in $O(|V| \times |E|)$ steps (the bound is optimal). Define $H_k(u, v)$ as the maximum weight of a path of length k from u to v in G , and fix an arbitrary source vertex $s \in V$. The algorithm proceeds by first iteratively computing the families $(H_k(s, v))_{v \in V}$ for increasingly large values of $k \leq |V|$ (each step has a cost $|E|$, by testing all the edges one by one to lengthen current maximal paths), and then returns $\beta = \max_{v \in V} \min_{0 \leq k < |V|} \frac{H_{|V|}(s, v) - H_k(s, v)}{|V| - k}$ (this final max min computation takes $O(|V|^2)$ steps, and as G is strongly connected, we know that $|E| \geq |V|$).

By giving the start weight $\varphi(u)$ to each edge (u, v) (or the end weight $\varphi(v)$, which doesn’t change the weight of cycles) on the De Bruijn graph G_r , using Karp’s algorithm computes β in $O(|\mathcal{A}|^{2r+1})$ steps of computation.

To describe $\mathcal{F}$, the main limitation of Karp’s algorithm is that it sort of stumbles upon β by happenstance, by finding (at least) one maximising cycle (in a not-so-obvious way, missed by Karp in his article offering a simpler but incorrect argument, and which took four decades to be noticed [CM17]), but without ever reaching a full knowledge of which vertices and cycles realise that maximum, and we could not find appropriate algorithms in the existing literature.

It might be enough to repeatedly run Karp’s algorithm for all potential sources $s \in V$ in order to identify all the relevant vertices, but this is not proven and rather than doing this we might as well adapt the algorithm itself directly. Hence, to compute F , taking inspiration from Karp’s algorithm, we propose the following procedure. First, compute $(H_k(u, v))_{u, v \in V}$ inductively for increasingly large values of $k \leq |V|$: we can initialise $H_0 = 0$, and given H_k , start with $h_{k+1} = -\infty$ and iterate over $u \in V$ and $e = (x, v) \in E$ to replace the current value of $h_{k+1}(u, v)$ by $H_k(u, x) + \varphi(x, v)$ if it is bigger, so that $h_{k+1} = H_{k+1}$ when the iteration

stops (*i.e.* computing H_{k+1} takes $O(|V| \times |E|)$ steps once H_k is known). Finally, simply have $F(u) = \max_{1 \leq k \leq |V|} \frac{H_k(u,u)}{k}$. This algorithm allows us to compute not only the mean max weight but also which vertices can realise it, in $O(|V|^2 |E|)$ steps of computation.

Once reframed in our setting (by shifting the weights from the vertices to the edges), we can compute the function $F : \mathcal{A}^r \rightarrow \mathbb{F}$ (and thus deduce $\mathcal{F}$ from the maximising set of F) in $O(|\mathcal{A}|^{3r+1})$ steps. This algorithm has been implemented and tested in both Python and Rust, and the code is available on a public repository [Gay26].

Remark 5.7. *As mentioned, Karp’s algorithm is optimal for arbitrary graphs. While this is beyond the scope of the paper, we conjecture here that for similar reasons, the $O(|V|^2 |E|)$ bound we obtain is optimal to produce all the vertices that maximise the mean cycle weight.*

However, in both cases, this doesn’t mean that there is no room for (algorithmic) improvements for our specific question, either because unlike in the aforementioned algorithms, our weights are on the vertices rather than the edges, or because there are specific optimisations to be found for the De Bruijn graphs.

References

- [AGS08] Luigi Ambrosio, Nicola Gigli and Giuseppe Savaré. *Gradient Flows in Metric Spaces and in the Space of Probability Measures*. 2nd ed. Birkhäuser Basel, 2008. [10.1007/978-3-7643-8722-8](#).
- [BGR12] Mark Braverman, Alexander Grigo and Cristóbal Rojas. *Noise vs Computational Intractability in Dynamics*. In: *ITCS*. ACM, 2012, pp. 128–141. [10.1145/2090236.2090247](#).
- [BGZ11] Jorge Buescu, Daniel Graça and Ning Zhong. *Computability and Dynamical Systems*. In: *Dynamics, Games and Science I*. Springer, 2011, pp. 169–181. [10.1007/978-3-642-11456-4_11](#).
- [Bou01] Thierry Bousch. *La Condition de Walters*. In: *Annales scientifiques de l’École Normale Supérieure* 34.2 (2001), pp. 287–311. [10.1016/s0012-9593\(00\)01062-4](#).
- [Boy+15] Laurent Boyer, Martin Delacourt, Victor Poupet, Mathieu Sablik and Guillaume Theyssier. *μ -Limit Sets of Cellular Automata from a Computational Complexity Perspective*. In: *Journal of Computer and System Sciences* 81.8 (2015), pp. 1623–1647. [10.1016/j.jcss.2015.05.004](#).
- [BP03] Vasco Brattka and Gero Presser. *Computability on Subsets of Metric Spaces*. In: *Theoretical Computer Science* 305 (2003), pp. 43–76. [10.1016/S0304-3975\(02\)00693-X](#).
- [Bré08] Julien Brémont. *Entropy and Maximizing Measures of Generic Continuous Functions*. In: *Comptes Rendus. Mathématique* 346.3–4 (2008), pp. 199–201. [10.1016/j.crma.2008.01.006](#).
- [Bur+22] Michael Burr, Suddhasattwa Das, Christian Wolf and Yun Yang. *Computability of Topological Pressure on Compact Shift Spaces Beyond Finite Type*. In: *Nonlinearity* 35.8 (2022), pp. 4250–4282. [10.1088/1361-6544/ac7702](#).

- [BW20] Michael Burr and Christian Wolf. *Computability at Zero Temperature*. In: *Nonlinearity* 33 (2020), pp. 6157–6175. [10.1088/1361-6544/ab9c71](https://doi.org/10.1088/1361-6544/ab9c71).
- [BW24] Michael Burr and Christian Wolf. *Computability in Dynamical Systems*. In: *Recent Developments in Fractal Geometry and Dynamical Systems*. AMS, 2024, pp. 85–98. [10.1090/conm/797](https://doi.org/10.1090/conm/797).
- [CM17] Mmanu Chaturvedi and Ross McConnell. *A Note on Finding Minimum Mean Cycle*. In: *Information Processing Letters* 127 (2017), pp. 21–22. [10.1016/j.ipl.2017.06.007](https://doi.org/10.1016/j.ipl.2017.06.007).
- [Col20] Pieter Collins. *Computable Analysis with Applications to Dynamic Systems*. In: *Mathematical Structures in Computer Science* 30.2 (2020), pp. 173–233. [10.1017/S096012952000002X](https://doi.org/10.1017/S096012952000002X).
- [Con16] Gonzalo Contreras. *Ground States are Generically a Periodic Orbit*. In: *Inventiones Mathematicae* 205 (2016), pp. 383–412. [10.1007/s00222-015-0638-0](https://doi.org/10.1007/s00222-015-0638-0).
- [CR24] Jordan Cotler and Semon Rezhikov. *Computational Dynamical Systems*. In: *IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*. 2024, pp. 166–202. [10.1109/FOCS61266.2024.00021](https://doi.org/10.1109/FOCS61266.2024.00021).
- [EM20] Aernout van Enter and Jacek Mi kisz. *Typical Ground States for Large Sets of Interactions*. In: *Journal of Statistical Physics* 181 (2020), pp. 1906–1914. [10.1007/s10955-020-02647-4](https://doi.org/10.1007/s10955-020-02647-4).
- [ENT23] Solène Esnay, Alonso Núñez and Ilkka Törmä. *Arithmetical Complexity of the Language of Generic Limit Sets of Cellular Automata*. In: *Journal of Computer and System Sciences* 134 (2023), pp. 20–41. [10.1016/j.jcss.2023.01.002](https://doi.org/10.1016/j.jcss.2023.01.002).
- [Gay26] Léo Gayral. *Complete Max Mean Weight Cycles*. PLMlab Repository. 2026. <https://plmlab.math.cnrs.fr/lgayral/complete-max-mean-weight-cycles>.
- [GHR11] Stefano Galatolo, Mathieu Hoyrup and Cristóbal Rojas. *Dynamics and Abstract Computability: Computing Invariant Measures*. In: *Discrete and Continuous Dynamical Systems* 29.1 (2011), pp. 193–212. [10.3934/dcds.2011.29.193](https://doi.org/10.3934/dcds.2011.29.193).
- [GST25] Léo Gayral, Mathieu Sablik and Siamak Taati. *Characterisation of the Set of Ground States of Uniformly Chaotic Finite-Range Lattice Models*. In: *Nonlinearity* 38.7 (June 2025). [10.1088/1361-6544/addbb8](https://doi.org/10.1088/1361-6544/addbb8).
- [GZ24] Daniel Graça and Ning Zhong. *Robust Non-Computability of Dynamical Systems and Computability of Robust Dynamical Systems*. In: *Logical Methods in Computer Science* 20.2 (2024). [10.46298/lmcs-20\(2:19\)2024](https://doi.org/10.46298/lmcs-20(2:19)2024).
- [Han74] William Hanf. *Nonrecursive Tilings of the Plane. I*. In: *Journal of Symbolic Logic* 39.2 (1974), pp. 283–285. [10.2307/2272640](https://doi.org/10.2307/2272640).
- [HM10] Michael Hochman and Tom Meyerovitch. *A Characterization of the Entropies of Multidimensional Shifts of Finite Type*. In: *Annals of Mathematics* 171.3 (2010), pp. 2011–2038. [10.4007/annals.2010.171.2011](https://doi.org/10.4007/annals.2010.171.2011).
- [HN26] Joel Hamkins and Theodor Nenu. *Did Turing Prove the Undecidability of the Halting Problem?* In: *Journal of Logic and Computation* 36.1 (Jan. 2026), exaf075. [10.1093/logcom/exaf075](https://doi.org/10.1093/logcom/exaf075).

- [Hoy07] Mathieu Hoyrup. *Dynamical Systems: Stability and Simulability*. In: *Mathematical Structures in Computer Science* 17.2 (2007), pp. 247–259. [10.1017/S096012950700597X](#).
- [HR09] Mathieu Hoyrup and Cristóbal Rojas. *Computability of Probability Measures and Martin-Löf Randomness Over Metric Spaces*. In: *Information and Computation* 207.7 (2009), pp. 830–847. [10.1016/j.ic.2008.12.009](#).
- [HR21] Mathieu Hoyrup and Jason Rute. *Computable Measure Theory and Algorithmic Randomness*. In: *Handbook of Computability and Complexity in Analysis*. Springer, 2021, pp. 227–270. [10.1007/978-3-030-59234-9_7](#).
- [HS18] Benjamin Hellouin de Menibus and Mathieu Sablik. *Characterization of Sets of Limit Measures of a Cellular Automaton Iterated on a Random Configuration*. In: *Ergodic Theory and Dynamical Systems* 38.2 (2018), pp. 601–650. [10.1017/etds.2016.46](#).
- [Jen06] Oliver Jenkinson. *Ergodic Optimization*. In: *Discrete and Continuous Dynamical Systems* 15.1 (2006). [10.3934/dcds.2006.15.197](#).
- [Jen19] Oliver Jenkinson. *Ergodic Optimization in Dynamical Systems*. In: *Ergodic Theory and Dynamical Systems* 39.10 (2019). [10.1017/etds.2017.142](#).
- [JV15] Emmanuel Jeandel and Pascal Vanier. *Characterizations of Periods of Multi-Dimensional Shifts*. In: *Ergodic Theory and Dynamical Systems* 35.2 (2015), pp. 431–460. [10.1017/etds.2013.60](#).
- [Kar78] Richard Karp. *A Characterization of the Minimum Cycle Mean in a Digraph*. In: *Discrete Mathematics* 23.3 (1978), pp. 309–311. [10.1016/0012-365X\(78\)90011-0](#).
- [KH03] Anatole Katok and Boris Hasselblatt. *A First Course in Dynamics: with a Panorama of Recent Developments*. Cambridge University Press, 2003. [10.1017/CB09780511998188](#).
- [LY75] Tien-Yien Li and James Yorke. *Period Three Implies Chaos*. In: *The American Mathematical Monthly* 82.10 (1975). [10.2307/2318254](#).
- [Mey11] Tom Meyerovitch. *Growth-Type Invariants for $\mathbb{Z}^d$ Subshifts of Finite Type and Arithmetical Classes of Real Numbers*. In: *Inventiones Mathematicae* 184.3 (2011), pp. 567–589. [10.1007/s00222-010-0296-1](#).
- [Mor10] Ian Morris. *Ergodic Optimization for Generic Continuous Functions*. In: *Discrete and Continuous Dynamical Systems* 27.1 (2010), pp. 383–388. [10.3934/dcds.2010.27.383](#).
- [Mye74] Dale Myers. *Nonrecursive Tilings of the Plane. II*. In: *Journal of Symbolic Logic* 39.2 (1974), pp. 286–294. [10.2307/2272641](#).
- [RY20] Cristóbal Rojas and Michael Yampolsky. *How to Lose at Monte Carlo: A Simple Dynamical System Whose Typical Statistical Behavior Is Non-Computable*. In: *Proceedings of the 52nd Symposium on Theory of Computing*. 2020, pp. 1066–1072. [10.1145/3357713.3384237](#).
- [RY24] Cristóbal Rojas and Michael Yampolsky. *Ulam Meets Turing: Constructing Quadratic Maps With Non-Computable SRB Measures*. 2024. [arXiv:2501.00006](#).

- [Sha64] Oleksandr Sharkovsky. *Coexistence of Cycles of a Continuous Map of the Line Into Itself*. In: *Ukrains'kyi Matematychnyi Zhurnal* 16.1 (1964), pp. 61–71.
- [Sha95] Oleksandr Sharkovsky. *Coexistence of Cycles of a Continuous Map of the Line Into Itself*. Trans. by Juan Tolosa. In: *International Journal of Bifurcation and Chaos* 5.5 (1995), pp. 1263–1273. [10.1142/S0218127495000934](#).
- [Spa08] Christoph Spandl. *Computability of Topological Pressure for Sofic Shifts With Applications in Statistical Physics*. In: *Journal of Universal Computer Science* 14.6 (2008), pp. 876–895. [10.3217/jucs-014-06-0876](#).
- [Tur37] Alan Turing. *On Computable Numbers, With an Application to the Entscheidungsproblem*. In: *Proceedings of the London Mathematical Society* s2-42 (1937), pp. 230–265. [10.1112/plms/s2-42.1.230](#).
- [Wei00] Klaus Weihrauch. *Computable Analysis*. Springer, 2000. [10.1007/978-3-642-56999-9](#).
- [Wes17] Linda Westrick. *Seas of Squares With Sizes From a Π_1^0 Set*. In: *Israel Journal of Mathematics* 222 (2017), pp. 431–462. [10.1007/s11856-017-1596-6](#).