On the typical periodic optimization (TPO) property

Bogdan Grechuk • 25 Mar 2026

In a paper recently accepted to Inventiones mathematicae (Huang et al. 2026), Huang, Jenkinson, Xu and Zhang proved that every S-gap shift has the typical periodic optimization (TPO) property.

Ergodic optimization asks a deceptively simple question: given a dynamical system (X, T) and an observable $f\colon X \to \mathbb{R}$, which invariant measures maximize the average value

$$\int f \, d\mu?$$

Even in very concrete systems, the set of invariant measures is huge and complicated, so one might expect maximizing measures to be difficult and irregular objects. A central theme of the subject is that, quite often, the opposite happens: for a "typical" regular observable, the maximizing measure is not exotic at all, but instead is supported on a single periodic orbit. This phenomenon is called *typical periodic optimization* (TPO), and it provides a striking bridge between two rather different viewpoints in dynamics: the global variational problem of maximizing an integral over all invariant measures, and the finite, combinatorial structure of periodic orbits.

To formulate this precisely, let (X, T) be a dynamical system, where $X = (X, d_X)$ is a compact metric space and $T\colon X \to X$ is continuous. Write $\mathcal{M}(X, T)$ for the set of all T-invariant Borel probability measures on X, that is,

$$\mathcal{M}(X, T) := \{\mu : \mu(T^{-1}(A)) = \mu(A) \text{ for every measurable } A \subseteq X\}.$$

A point $x \in X$ is *periodic* if $T^n(x) = x$ for some $n \geq 1$. If x is periodic of period n, then the probability measure

$$\mu_x := \frac{1}{n} \sum_{k=0}^{n-1} \delta_{T^k x}$$

is T-invariant; such measures are called *periodic measures*. For optimization problems, these are the simplest invariant measures one could hope for, since they are supported on finite orbits and can therefore be described explicitly.

We will work with Lipschitz observables. A function $f\colon X \to \mathbb{R}$ is *Lipschitz* if there exists $L \geq 0$ such that

$$|f(x) - f(y)| \leq L\, d_X(x, y) \qquad \text{for all } x, y \in X.$$

The space $\mathrm{Lip}(X)$ becomes a normed space when equipped with

$$\|f\|_{\mathrm{Lip}} := \sup_{x \in X} |f(x)| + \sup_{x \neq y} \frac{|f(x) - f(y)|}{d_X(x, y)}.$$

This is a natural class in ergodic optimization: it is large enough to include many observables of interest, but regular enough that one can hope to exploit the geometry of the dynamics.

A continuous function $f\colon X \to \mathbb{R}$ is said to have the *periodic optimization property* if there exists a periodic measure $\mu_x \in \mathcal{M}(X, T)$ such that

$$\int f\, d\mu_x > \int f\, d\mu \qquad \text{for every } \mu \in \mathcal{M}(X, T) \setminus \{\mu_x\}.$$

In other words, f has a unique maximizing measure, and that maximizing measure is periodic. The system (X, T) is said to have the *typical periodic optimization* property if there exists an open dense set $U \subseteq \mathrm{Lip}(X)$ such that every $f \in U$ has the periodic optimization property.

This definition is worth pausing over. Openness says that the maximizing periodic orbit is stable under small perturbations of the observable, so the phenomenon is robust rather than accidental. Density says that periodic optimization is not a rare curiosity but is arbitrarily close to any Lipschitz observable. Thus TPO gives a very strong structural description of generic optimization problems on the system. It says that, although the simplex $\mathcal{M}(X, T)$ may be enormously complicated, the generic maximization problem is governed by a finite orbit. From both a conceptual and computational point of view, this is extremely useful.

At first sight, TPO is rather surprising. Periodic measures form a very small subset of $\mathcal{M}(X, T)$, and many dynamical systems support highly non-periodic invariant measures with rich statistical behaviour. One might therefore expect generic maximizing measures to be non-periodic as well. Nevertheless, numerical and heuristic evidence from the 1990s suggested that periodic maximizers often dominate in concrete optimization problems. The challenge was to explain this rigorously and to identify broad classes of systems for which it holds.

A decisive advance was made by Contreras (Contreras 2016), who proved TPO for expanding maps. Recall that a continuous map $T\colon X \to X$ is *expanding* if it is Lipschitz and there exist $d \in \mathbb{Z}^+$ and $0 < \lambda < 1$ such that for every $x \in X$ there is a neighbourhood U_x together with continuous inverse branches S_i $(i = 1, \ldots, l_x \leq d)$ having pairwise disjoint images and satisfying

$$T^{-1}(U_x) = \bigcup_{i=1}^{l_x} S_i(U_x), \qquad T \circ S_i = \mathrm{Id}_{U_x},$$

and

$$d_X(S_i(y), S_i(z)) \le \lambda\, d_X(y,z) \qquad \text{for all } y, z \in U_x.$$

This formalizes the idea that inverse iterates contract distances uniformly. Such systems are strongly hyperbolic, and periodic orbits are abundant and well behaved, so one can hope to perturb an observable in a controlled way until one periodic orbit becomes strictly dominant.

Theorem 1 *If X is a compact metric space and $T\colon X \to X$ is an expanding map, then (X,T) has the TPO property.*

The significance of this theorem goes beyond the expanding setting itself. It showed that TPO is a genuine and robust dynamical phenomenon, not merely an artefact of examples. At the same time, it left open an important question: how much hyperbolicity is really needed? Many natural systems are not expanding, at least not globally, and symbolic systems in particular often display a mixture of rigid combinatorial constraints and non-expanding behaviour. Extending TPO beyond the uniformly expanding world therefore became a major problem.

A recent breakthrough in this direction was obtained by Huang, Jenkinson, Xu and Zhang (Huang et al. 2026), who developed a general framework proving TPO for broad classes of non-expanding systems. One especially attractive family of examples is given by S-gap shifts. These are symbolic systems built from a very simple rule—the allowed numbers of zeros between consecutive ones—yet they exhibit a rich range of behaviours, from shifts of finite type to non-sofic systems. They therefore provide an excellent testing ground for the scope of modern ergodic optimization methods.

Let

$$\Sigma := \{0,1\}^{\mathbb{N}}$$

be the full one-sided shift space, equipped with the metric d defined by $d(x,x) = 0$ and

$$d(x,y) := 2^{-N(x,y)} \qquad \text{for } x \neq y,$$

where

$$N(x,y) := \min\{n \ge 0 : x_{n+1} \neq y_{n+1}\}.$$

Thus two sequences are close precisely when they agree for a long initial block. The left shift

$$\sigma(x_1, x_2, x_3, \dots) = (x_2, x_3, x_4, \dots)$$

is continuous on Σ.

Now fix a non-empty set $S \subseteq \mathbb{N} \cup \{0\}$. If $k \geq 0$, write 0^k for the word consisting of k consecutive zeros, with 0^0 understood as the empty word. Define

$$E_S := \{10^{s_1}10^{s_2}10^{s_3} \cdots \in \Sigma : \; s_i \in S \text{ for all } i \geq 1\},$$

and let

$$X_S := \overline{\bigcup_{n=0}^{\infty} \sigma^n(E_S)} \subseteq \Sigma.$$

The shift space X_S is called the *S-gap shift*. Informally, its points are precisely those binary sequences in which the gaps between successive 1's have lengths prescribed by S, together with all sequences forced by taking shifts and limits. This definition is simple enough to keep concrete examples in view, but flexible enough to produce highly nontrivial symbolic dynamics.

The interest of S-gap shifts lies partly in this balance between simplicity and complexity. The entire system is controlled by the arithmetic set S, yet small changes in S can produce markedly different dynamical behaviour. Some S-gap shifts are shifts of finite type, some are sofic, and many are neither. From the perspective of ergodic optimization, they are therefore a natural class in which to ask whether periodic optimization remains typical once one moves beyond classical uniformly expanding models.

The answer is yes.

Theorem 2 *Every S-gap shift (X_S, σ) has the typical periodic optimization property.*

This theorem is compelling for several reasons. First, it shows that TPO persists in a setting where one no longer has the straightforward geometric expansion present in Theorem 1. Second, it applies uniformly to *every* choice of S, so the conclusion is not tied to special arithmetic assumptions on the allowed gaps. Third, the result fits into a broader picture in which symbolic structure, rather than metric expansion alone, can force generic maximizing measures to be periodic.

From a broader viewpoint, Theorem 2 reinforces the idea that periodic orbits are not merely useful approximations to maximizing behaviour: in many systems, they are the generic optimizers. This is conceptually satisfying, because periodic points are the most concrete invariant objects in the system, and it is also practically valuable, since optimization over periodic orbits is often far more tractable than optimization over all invariant measures. In this sense, TPO converts an infinite-dimensional variational problem into a finite-dimensional one for a generic observable.

Theorem 2 is only one instance of a much broader theory developed in (Huang et al. 2026). The real achievement of that work is not just the treatment of S-gap shifts, but the emergence of a flexible mechanism for proving TPO in symbolic systems that are far from uniformly expanding. Theorem 2 is therefore best viewed both as a striking result in its own right and as evidence for a larger principle: generic optimization in dynamics is often governed by periodic structure even in settings where that structure is far from obvious at the outset.

References

Contreras, Gonzalo. 2016. "Ground States Are Generically a Periodic Orbit." *Invent. Math.* 205 (2): 383–412. https://doi.org/10.1007/s00222-015-0638-0.

Huang, Wen, Oliver Jenkinson, Leiye Xu, and Yiwei Zhang. 2026. "Typical Periodic Optimization for Dynamical Systems: Symbolic Dynamics." *Inventiones Mathematicae*, ahead of print. https://doi.org/10.1007/s00222-026-01411-x.