

Descriptive complexity of topological invariants

Djamel Eddine Amir, Mathieu Hoyrup

▶ To cite this version:

Djamel Eddine Amir, Mathieu Hoyrup. Descriptive complexity of topological invariants. 2023. hal-04222391

HAL Id: hal-04222391 https://inria.hal.science/hal-04222391

Preprint submitted on 29 Sep 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Descriptive complexity of topological invariants

Djamel Eddine Amir and Mathieu Hoyrup Université de Lorraine, CNRS, Inria, LORIA, Nancy, 54000, France

September 29, 2023

Abstract

In this article, we investigate the descriptive complexity of topological invariants. Our main goal is to understand the expressive power of low complexity invariants, by investigating which spaces they can distinguish. We study the invariants in the first two levels of the Borel hierarchy. We develop techniques to establish whether two spaces can be separated by invariants in these levels. We show that they are sufficient to separate finite topological graphs. We finally identify the complexity of recognizing the line segment.

1 Introduction

Given a subset of the plane or of a Euclidean space, described in some way, how difficult is it to analyze the topology of that set? How difficult is it to recognize a particular space, i.e. to test that the given set is homeomorphic to that space? How difficult is it to distinguish two particular spaces? More generally, how difficult is it to test whether the set satisfies a given topological invariant?

Those questions can be investigated in many different ways, depending on the class of spaces that are allowed, the way they are described, and on the notion of complexity which is used to measure the difficulty of detecting topological invariants.

In this article, we work with compact subsets of Euclidean spaces, and more generally compact subsets of the Hilbert cube $[0,1]^{\mathbb{N}}$. The description of such a subset essentially tells us which voxels or rational balls intersect it. This description is weak, in the sense that it does not carry any structural or combinatorial information; for instance, if the space is a topological graph, it does not give direct access to the combinatorial description of the graph. Therefore, no topological invariant can be decided in finite time, so the relevant measure of complexity is *descriptive* complexity. It measures the difficulty of expressing a property in terms of simpler properties, and in its effective version, the difficulty of expressing the property by a logical formula, or by a program.

The main purpose of topological invariants is to separate different topological spaces, i.e. to show that they are not homeomorphic. A topological invariant is interesting if it is able to tell different spaces apart, and if it leads to a rich theory. From a logical or computational viewpoint, it is also interesting if it is simple to test.

Our general goal in this article is to understand what topological invariants can be tested or described using a limited amount of complexity. We study the following particular problems:

- As far as possible, characterize the topological invariants of a given complexity level,
- For a restricted class of spaces, in particular for a pair of spaces, identify the minimal level of complexity that allows to distinguish all these spaces,
- Given a particular space, find the complexity of recognizing that space.

The descriptive complexity of topological invariants has been previously studied. The usual approach is to identify the descriptive complexity of pre-existing invariants, provided by topology. In our approach we go the other way round, starting from a complexity level and looking for expressive invariants having this level of complexity, by investigating which spaces can be separated using invariants of that complexity. We focus in particular on invariants of low complexity, namely Π_1^0 and Σ_2^0 . We mention that Σ_2^0 invariants have recently proved useful in studying the computability properties of topological spaces [AH22].

Let us mention previous works on the descriptive complexity of topological invariants. Ajtaj and Becker proved that path-connectedness is Π_2^1 -complete. Becker proved that for compact subsets of \mathbb{R}^2 , simply-connectedness is Π_1^1 -complete ([Bec92, Kec95]). Debs and Saint-Raymond proved that connectedness is Π_1^0 , local connectedness if Π_3^0 , and gave another, independent proof that path-connectedness is Π_2^1 [DS20]. Lupini, Melnikov and Nies [LMN23] and Downey, Melnikov [DM23] showed that the Čech cohomology groups are computable; in particular, whether the *n*th Čech cohomology group $\check{H}^n(X)$ is non-trivial is Σ_2^0 . The general theory of the descriptive complexity of invariants was studied by Vaught [Vau74].

When more information on the space is available, for instance if the space is given by a finite triangulation, many topological invariants become decidable, and the appropriate notion of complexity is their *computational* complexity, typically their time complexity. For instance, the complexity of combinatorial invariants associated to 3-manifolds represented by triangulations was studied by Burton, Maria and Spreer in [BMS15]. For simplicial complexes, the probem of deciding whether a homology group is non-trivial was proved to be NP-hard by Adamaszek and Stacho in [AS16]. The complexity of certain invariants associated to simply-connected spaces represented in an algebraic way was studied by Amann in [Ama15].

1.1 Overview of the results

In this section we summarize the main results of the article. We start explaining briefly how the descriptive complexity of topological invariants is measured (more details in Section 1.2).

Descriptive complexity of topological invariants. Every compact metrizable space embeds in the Hilbert cube $Q = [0, 1]^{\mathbb{N}}$ endowed with the product topology. Therefore, such a space can be described as a compact subset of Q. The space $\mathcal{K}(Q)$ of non-empty compact subsets of Q can in turn be endowed with the Vietoris topology, or equivalently the Hausdorff metric.

A property of compact sets is a set $\mathcal{P} \subseteq \mathcal{K}(Q)$. It is Π_1^0 if it is closed in the Vietoris topology. It is Σ_2^0 if it is a countable union of Π_1^0 properties. Effective complexity classes Π_1^0 and Σ_2^0 can be defined, by requiring that the corresponding description of the set can be produced by a program. A **topological invariant** is a property $\mathcal{P} \subseteq \mathcal{K}(Q)$ such that if $X, Y \in \mathcal{K}(Q)$ are homeomorphic, then $X \in \mathcal{P}$ iff $Y \in \mathcal{P}$.

Characterization of the Π_1^0 topological invariants (Section 2). We first give a complete characterization of the invariants of complexity Π_1^0 , showing that they can only express connectedness properties of the space. Moreover, the non-effective class of Π_1^0 invariants and the effective class of Π_1^0 invariants actually coincide.

For $0 \leq p \leq n$, we define an invariant $\mathcal{C}_{n,p}$ as follows: $X \in \mathcal{C}_{n,p}$ iff X has at most n connected components, among which at most p are not singletons.

Theorem (Theorem 2.1). Let \mathcal{P} be a non-trivial topological invariant. The following statements are equivalent:

- 1. $\mathcal{P} \in \Pi_1^0$,
- 2. $\mathcal{P} \in \Pi_1^0$,
- 3. \mathcal{P} is a finite union of $\mathcal{C}_{n,p}$'s.

The Σ_2^0 invariants (Section 3). We carry out a thorough study of the class of Σ_2^0 invariants. We obtain a characterization of when a space can be separated from another by some Σ_2^0 invariant. We say that a compact space X strongly approximates a compact space Y if there is a copy $X_0 \subseteq Q$ of X such that for every $\epsilon > 0$, one can continuously deform X_0 to converge to a copy of Y, in such a way that every point of X_0 is moved by at most ϵ .

Theorem (Theorem 3.1). Let X, Y be compact spaces. The following statements are equivalent:

- X strongly approximates Y,
- Every \sum_{2}^{0} invariant satisfied by X is satisfied by Y.

We apply this result to show that the line segment and the closed topologist's sine curve cannot be separated by Σ_2^0 invariants, and give other applications.

Separating finite topological graphs (Section 4). Separating arbitrary spaces may require topological invariants of arbitrarily high complexity. However, when restricting to certain families of spaces, one might expect that some level of complexity is sufficient to separate all these spaces, raising the problem of finding the optimal level of complexity needed.

We investigate this problem for the family of finite topological graphs. We first show that the non-effective class Σ_2^0 is no more expressive than the effective class Σ_2^0 .

Theorem (Theorem 4.1). Let X be a finite topological connected graph and Y a compact space. If there exists a Σ_2^0 invariant satisfied by X but not by Y, then there exists a Σ_2^0 such invariant.

Next, we show that the Σ_2^0 invariants are expressive enough to separate finite graphs.

Theorem (Theorem 4.3). If G, H are non-homeomorphic finite topological graphs, then some Σ_2^0 invariant is satisfied by G but not by H, or vice-versa.

Note that there is a dissymmetry: it can happen that every Σ_2^0 invariant satisfied by G is also satisfied by H (but not vice-versa). We give a characterization of this relation. A graph H can be contracted to a graph G if there is a sequence of edge contractions from H to G; an edge contraction consists in removing an edge and merging its two endpoints.

Theorem (Theorem 4.2). The following statements are equivalent for finite connected graphs G, H:

- Every Σ_2^0 invariant satisfied by G is satisfied by H,
- Some subdivision of H can be contracted to G.

Recognizing the line segment (Section 5). To any compact space X is associated the following problem of recognizing X:

Input. A compact set Y,

Output. Is Y homeomorphic to X?

Note that this problem is a topological invariant. A classical result from Descriptive Set Theory implies that for each particular X, the descriptive complexity of recognizing X is Borel. It suggests the following research program: identify, for each concrete space X, the exact complexity of recognizing X. We make a first step in this direction, solving this question when X is the line segment.

Theorem (Theorem 5.1). The problem of recognizing the line segment is Π_3^0 -complete.

1.2 Background

We give definitions of the main notions that are needed in this article.

Descriptive complexity. The Hilbert cube $Q = [0, 1]^{\mathbb{N}}$ is endowed with the product topology, induced by the metric $d(x, y) = \sum_{i \in \mathbb{N}} 2^{-i} |x_i - y_i|$, where $x = (x_0, x_1, \ldots)$ and $y = (y_0, y_1, \ldots)$. A point $x \in Q$ is rational if its coordinates are rational and it has finitely many non-zero coordinates. The topology is generated by the rational balls centered at rational points, with positive rational radii.

We consider the space $\mathcal{K}(Q)$ of non-empty compact subsets of Q, endowed with the **Vietoris topology**, generated by the sets $\{X \in \mathcal{K}(Q) : X \cap U \neq \emptyset\}$ and $\{X \in \mathcal{K}(Q) : X \subseteq U\}$, where U ranges over the open subsets of Q. A basis for the Vietoris topology is obtained by taking finite intersections of these sets, where U is a finite union of rational balls. This topology is also induced by the Hausdorff metric d_H defined by $d_H(X,Y) = \max(\max_{x \in X} \min_{y \in Y} d(x, y), \max_{y \in Y} \min_{x \in X} d(x, y)).$

We will use the fact that points of Q can be multiplied by scalars in [0, 1], and that convex combinations can be performed in Q.

A property $\mathcal{P} \subseteq \mathcal{K}(Q)$ is Σ_1^0 if it is open in the Vietoris topology. Inductively on $n \in \mathbb{N}$, a property $\mathcal{P} \subseteq \mathcal{K}(Q)$ is Π_n^0 if its complement is Σ_n^0 , and \mathcal{P} is Σ_{n+1}^0 if $\mathcal{P} = \bigcup_n \mathcal{P}_n$, where \mathcal{P}_n are Π_n^0 .

The effective classes are defined as follows. $\mathcal{P} \subseteq \mathcal{K}(Q)$ is Σ_1^0 if there exists a program enumerating an infinite list of basic Vietoris open sets (encoded in a standard way) whose union is \mathcal{P} . \mathcal{P} is Π_n^0 if its complement is Σ_n^0 , and \mathcal{P} is Σ_{n+1}^0 if $\mathcal{P} = \bigcup_n \mathcal{P}_n$ where each \mathcal{P}_n is Π_n^0 , uniformly in n (i.e. there is a single program describing \mathcal{P}_n on input n for all n).

We recall that a property $\mathcal{P} \subseteq \mathcal{K}(Q)$ is a topological invariant if for $X, Y \in \mathcal{K}(Q)$, if $X \in \mathcal{P}$ and Y is homeomorphic to X, then $Y \in \mathcal{P}$.

In this article, all the topological spaces are compact metrizable spaces, which we may implicitly assume. A **copy** of such a space X is a subset of Q that is homeomorphic to X.

2 Π_1^0 invariants

In this section, we give a complete characterization of the topological invariants in the lowest descriptive complexity class, namely $\mathbf{\Pi}_{1}^{0}$. We show that this level of complexity can

only detect connectedness properties of the space.

In a topological space X, a subset A is **connected** if there is no disjoint open sets $U, V \subseteq X$ both intersecting A and such that $A \subseteq U \cup V$. A **connected component** of X is a maximal connected subset. We say that it is trivial if it is a singleton.

Definition 2.1. Let $0 \leq p \leq n$ be natural numbers. The topological invariant $C_{n,p}$ is defined by: $X \in C_{n,p}$ iff X has at most n connected components, among which at most p are non-trivial.

As particular cases,

- $X \in \mathcal{C}_{n,n} \iff X$ has at most *n* connected components,
- $X \in \mathcal{C}_{n,0} \iff X$ contains at most n points.

These invariants are examples of Π^0_1 invariants.

Proposition 2.1. Let $0 \le p \le n$. The invariant $C_{n,p}$ is Π_1^0 .

Proof. We show that $X \notin C_{n,p}$ iff there exist n + 1 disjoint open sets covering X, each one intersecting X, or p + 1 disjoint open sets covering X, each one intersecting X in at least n - p + 1 distinct points. As X is compact, one can quantify over finite unions of rational balls rather than general open sets, so this property is Σ_1^0 .

Of course, if $X \notin C_{n,p}$ then one of these coverings exists: either X has at least n + 1 connected components, or X has at least p + 1 non-trivial connected components, which then contain as many points as needed. Conversely, if the first type of covering exists, then X has more than n components so $X \notin C_{n,p}$. Now assume that the second type of covering exists. If X has more than p non-trivial components, then $X \notin C_{n,p}$. If X has at most p non-trivial components, then one of the open sets does not contain an infinite component, so it contains at least n - p + 1 singletons. In total, X contains at least p + (n - p + 1) = n + 1 components, so $X \notin C_{n,p}$.

Moreover, these invariants generate all the Π_1^0 invariants.

Theorem 2.1. Let \mathcal{P} be a non-trivial topological invariant. The following statements are equivalent:

- 1. $\mathcal{P} \in \Pi^0_1$,
- 2. $\mathcal{P} \in \Pi^0_1$,
- 3. \mathcal{P} is a finite union of $\mathcal{C}_{n,p}$'s.

Observe that $3. \Rightarrow 2$. follows from Proposition 2.1 and $2. \Rightarrow 1$. is obvious. The next section is devoted to the proof of $1. \Rightarrow 3$.

2.1 Approximations

We introduce the notion of approximation, which is the key ingredient underlying the proof.

Definition 2.2. A compact space X **approximates** Y if some copy $Y_0 \subseteq Q$ of Y is a limit of copies of X (in the Vietoris topology, or equivalently the Hausdorff metric).

Proposition 2.2. Let X, Y be compact spaces. The following statements are equivalent:

- X approximates Y,
- Every $\prod_{i=1}^{0}$ invariant satisfied by X is satisfied by Y.

Proof. If X approximates Y and X satisfies a Π_1^0 invariant \mathcal{P} , then \mathcal{P} contains all the copies of X as well as their limits, as \mathcal{P} is closed. Therefore, \mathcal{P} contains a copy of Y.

Conversely, assume that every Π_1^0 invariant satisfied by X is satisfied by Y. Let \mathcal{P} be the set of limits of copies of X. By definition, \mathcal{P} is closed, i.e. Π_1^0 . \mathcal{P} is a topological invariant: if $K \subseteq Q$ is a limit of copies X_n of X and $K' \subseteq Q$ is homeomorphic to K, then K' is a limit of copies of X, because the homeomorphism $f: K \to K'$ can be extended to a homeomorphism $F: Q \to Q$ (Theorem 5.2.4 in [vM01]), and the copies $F(X_n)$ of X converge to F(K) = K'. Strictly speaking, Theorem 5.2.4 in [vM01] can be applied only if $K, K' \subseteq (0, 1)^{\mathbb{N}}$. If it is not the case, we apply the argument to copies of K, K' scaled down so that they fit in $(0, 1)^{\mathbb{N}}$, and observe that K' is a limit of such copies, which are limits of copies of X.

Finally, \mathcal{P} is satisfied by X hence by Y, i.e. some copy of Y is a limit of copies of X. \Box

The proof also shows that in the definition of approximation, one could equivalently replace "some copy" by "every copy". The next result gives a particularly simple criterion to establish that a space approximates another space.

Proposition 2.3. Let X, Y be compact spaces. If there exists a continuous surjective function $f : X \to Y$, then X approximates Y.

Proof. We assume that X, Y are embedded in Q. The function $f : X \to Y$ is a limit of injective continuous functions $f_n : X \to Q$, defined as follows. We recall that elements of Q are sequences of real numbers. For $x \in Q$, let u be the prefix of f(x) of length n and let $f_n(x)$ be the concatenation of u and x. By definition, f_n is injective, continuous, and is closer and closer to f as n grows. Therefore, Y is the limit of the sets $f_n(X)$, which are copies of X.

If X, Y are topological spaces, then their **disjoint union** is the space $X \sqcup Y$, which can be defined as follows: if X and Y are embedded in Q, then $(\{0\} \times X) \cup (\{1\} \times Y)$ is a copy of $X \sqcup Y$ in $[0, 1] \times Q \cong Q$.

Proposition 2.4. 1. The approximation relation is a pre-order, i.e. it is reflexive and transitive,

- 2. If X_i approximates Y_i for $i \in \{0, 1\}$, then $X_0 \sqcup X_1$ approximates $Y_0 \sqcup Y_1$,
- 3. If $X \neq \emptyset$, then $X \sqcup Y$ approximates X,
- 4. If X, Y are connected and X is not a singleton, then X approximates Y.

Proof. We only prove 4., the others are easy.

As X is not a singleton, it contains two distinct points, which differ in at least one coordinate. The projection along that coordinate is a continuous function sending X to a line segment, so X approximates the line segment by Proposition 2.3. The line segment approximates every finite connected graph, because such a graph is a continuous image of the line segment (a path visiting all the edges, possibly with repetitions, is nothing else than a continuous surjective function from the line segment to the graph). It remains to show that every connected compact space Y is a limit of finite connected graphs.

Given $\epsilon > 0$, let $V \subseteq Y$ be a finite set such that every point of Y is ϵ -close to a point of V. Consider the graph G with vertex set V, with an edge between u and v if $d(u, v) < 2\epsilon$. G is naturally embedded in Q, and each point of an edge is ϵ -close to one of the endpoints, so $d_H(G, Y) < \epsilon$. G must be connected: if there is a non-trivial partition $V = V_1 \sqcup V_2$ with no edge between V_1 and V_2 , then the open ϵ -neighborhoods of V_1 and V_2 cover Y and do not intersect, so they disconnect Y, giving a contradiction. \Box

We now have all the ingredients to prove the main result of this section.

Proof of Theorem 2.1. Let \mathcal{P} be a non-trivial Π_1^0 invariant. We first show that there exists $k \in \mathbb{N}$ such that every space in \mathcal{P} has at most k connected components. We assume the contrary and show that every space belongs to \mathcal{P} , contradicting the non-triviality of \mathcal{P} . First, a space X with at least n components approximates the finite set F with n points, because there is a continuous surjective function $f: X \to F$. As the finite sets are dense in $\mathcal{K}(Q)$, if the closed set \mathcal{P} contains every finite set, it contains every compact set. Therefore, \mathcal{P} is trivial, giving the desired contradiction.

Let $X \in \mathcal{P}$ have *n* connected components, *p* of them being non-trivial. *X* approximates every set $Y \in \mathcal{C}_{n,p}$: to each connected component Y_i of *Y* one can associate a connected component X_i of *X* in an injective way, such that if Y_i is non-trivial then X_i is non-trivial. Each Y_i is approximated by X_i by Proposition 2.4, item 4., so *X* approximates *Y* by Proposition 2.4, items 2. and 3.. As a result, every $Y \in \mathcal{C}_{n,p}$ satisfies \mathcal{P} , i.e. $\mathcal{C}_{n,p} \subseteq \mathcal{P}$. Therefore, \mathcal{P} contains the union of the $\mathcal{C}_{n,p}$'s such that there exists $X \in \mathcal{P} \cap \mathcal{C}_{n,p}$, and is contained in that union because every $X \in \mathcal{P}$ has a finite number of connected components. Note that the union is finite because the *n*'s are upper bounded by *k*.

3 $\sum_{i=1}^{0} \sum_{j=1}^{0} invariants$

Our goal is to analyze the separation power of Σ_2^0 invariants. A natural invariant of that complexity is given by the topological dimension. A topological space X has dimension

at most n if for every open cover $(U_i)_{i\in\mathbb{N}}$ of X, there exists an open cover $(V_i)_{i\in\mathbb{N}}$ of X such that each V_i is contained in some U_j and at most n + 1 sets V_i 's have non-empty intersection.

Proposition 3.1. Having dimension at least n is a Σ_2^0 invariant.

Proof. Let X be a compact subset of Q. Using standard compactness arguments, one can show that $\dim(X) \ge n \iff$ there exists a finite sequence $U_0 \ldots, U_k$ of finite union of rational balls of Q covering X such that there is no finite sequence V_0, \ldots, V_l of finite unions of rational balls covering X, such that each closure \overline{V}_i is contained in some U_j and each intersection of n + 1 sets \overline{V}_i 's is empty. We now use the fact that Q is effectively compact, i.e. given a Σ_1^0 subset U of Q, one can semidecide whether U = Q (see [IK21] for instance); in other words, given a Π_1^0 subset of Q, one can semidecide whether it is empty. Therefore, whether a closed ball is contained in an open set, and whether a closed set is empty are semidecidable or Σ_1^0 conditions, so the whole formula is Σ_2^0 .

3.1 Strong approximation

We now introduce a refinement of the notion of approximation and prove an analog to Proposition 2.2.

Definition 3.1. Let $X \subseteq Q$ and $\epsilon > 0$. An ϵ -function is a continuous function $f : X \to Q$ such that $d(f(x), x) < \epsilon$ for all $x \in X$. An ϵ -deformation of X is the image of X under an ϵ -function.

Definition 3.2. Let X, Y be compact spaces. We say that X strongly approximates Y, written $X \leq Y$, if there exists a copy $X_0 \subseteq Q$ of X such that for every $\epsilon > 0$, some copy of Y is a limit of ϵ -deformations of X_0 .

Note that the definition would be equivalent if one requires the ϵ -deformations of X_0 to be images of X_0 under *injective* ϵ -functions. Indeed, as in Proposition 2.3, an ϵ -function is a limit of injective ϵ -functions. As for the notion of approximation, if X strongly approximates Y, then the condition actually holds for *every* copy $X_0 \subseteq Q$ of X.

Theorem 3.1. Let X, Y be compact spaces. The following statements are equivalent:

- X strongly approximates Y,
- Every $\sum_{n=1}^{\infty} 2^n$ invariant satisfied by X is satisfied by Y.

Proof. Assume that every Σ_2^0 invariant satisfied by X is satisfied by Y. Let $X_0 \subseteq Q$ be a copy of X and let $\epsilon > 0$. Consider the following topological invariant \mathcal{P} : $Y \in \mathcal{P}$ iff there exists $\delta < \epsilon$ and a copy Y_0 of Y which is a limit of δ -deformations of X_0 . We show that it is Σ_2^0 . Let \mathcal{P}_{δ} be the set of limits of δ -deformations of X_0 . It is closed, i.e. Π_1^0 . Let $(f_i)_{i \in \mathbb{N}}$ be a sequence of injective continuous functions $f_i : Q \to Q$ that is dense in the space of continuous functions from Q to itself with the metric d (it exists as the space of continuous functions is separable, by Proposition 1.3.3 in [vM01]).

Claim 3.1. For $Y \subseteq Q$, one has $Y \in \mathcal{P}$ if and only if there exists a positive rational $\delta < \epsilon$ and *i* such that $f_i(Y) \in \mathcal{P}_{\delta}$.

Proof. Of course, if δ and i exist, then $Y \in \mathcal{P}$ because $f_i(Y)$ is a copy of Y which is a limit of δ -deformations of X_0 . Conversely, assume that $Y \in \mathcal{P}$ and let $\delta < \epsilon$ and $g_n : X_0 \to Q$ be δ -functions such that $g_n(X_0)$ converge to a copy $Y_0 \subseteq Q$ of Y. Let $f: Y \to Y_0$ be a homeomorphism and $h: Q \to Q$ be a continuous extension of f^{-1} . Let $\mu < \epsilon - \delta$ and let f_i be μ -close to f for the metric d. One has $f_i(Y) = f_i \circ f^{-1}(Y_0) = f_i \circ h(Y_0) =$ $\lim_n f_i \circ h \circ g_n(X_0)$. Note that the restriction of $f_i \circ h$ to Y_0 is a μ -function, so its restriction to a neighborhood of Y_0 is a μ -function as well, by compactness of Y_0 . For sufficiently large $n, g_n(X_0)$ is contained in that neighborhood, so $f_i \circ h \circ g_n$ is a $(\mu + \delta)$ -function. Therefore, for any rational number δ' between $\mu + \delta$ and $\epsilon, f_i(Y)$ is a limit of δ' -deformations of X_0 , i.e. $f_i(Y) \in \mathcal{P}_{\delta'}$.

The condition in the claim is Σ_2^0 , so \mathcal{P} is Σ_2^0 . Observe that X obviously satisfies \mathcal{P} , so Y satisfies \mathcal{P} as well by assumption. As it is true for every $\epsilon > 0$, X strongly approximates Y.

Conversely, assume that X strongly approximates Y and let \mathcal{P} be a $\sum_{i=1}^{0} \mathbb{P}_{i}$ invariant satisfied by X. We fix some arbitrary copy $X_0 \subseteq Q$ of X. One has $\mathcal{P} = \bigcup_{n \in \mathbb{N}} \mathcal{P}_n$ where each \mathcal{P}_n is Π_1^0 . For every homeomorphism $f: Q \to Q$, $f(X_0) \in \mathcal{P}$ so $f(X_0) \in \mathcal{P}_n$ for some n. Therefore, there exists $n \in \mathbb{N}$ such that for "many" f, $f(X_0) \in \mathcal{P}_n$. It is made precise by using the Baire category theorem as follows. Let $\mathcal{H}(Q)$ be the topological space of homeomorphisms from Q to itself, with the topology induced by the metric d. It is a Polish space by Proposition 1.3.12 in [vM01] (note that d itself is not complete, but there exists another complete metric inducing the same topology). For each $n \in \mathbb{N}$, let $A_n = \{f \in \mathcal{H}(Q) : f(X_0) \in \mathcal{P}_n\}$. As \mathcal{P}_n is closed, A_n is closed as well. One has $\mathcal{H}(Q) = \bigcup_{n \in \mathbb{N}} A_n$ so by the Baire category theorem, there exists n such that A_n has non-empty interior, i.e. contains a ball $B(f, \epsilon)$. As X strongly approximates Y, there exist injective ϵ -functions $f_i : f(X_0) \to Q$ such that $f_i(f(X_0))$ converge to a copy Y_0 of Y. Note that each $f_i \circ f$ belongs to $B(f, \epsilon)$ which is contained in A_n , so $f_i(f(X_0)) \in \mathcal{P}_n$. As \mathcal{P}_n is closed and $f_i(f(X_0))$ converge to $Y_0, Y_0 \in \mathcal{P}_n$. Therefore, Y satisfies \mathcal{P} and the proof is complete.

3.1.1 The closed topologist's sine curve and the line segment

Theorem 3.1 gives a very effective way of proving that there is no Σ_2^0 invariant separating X from Y, i.e. satisfied by X but not by Y: one simply needs to design ϵ -deformations of X converging to a copy of Y, for every ϵ . We illustrate this technique by showing that the line segment and the closed topologist's sine curve cannot be separated by Σ_2^0 invariants.

Definition 3.3. The closed topologist's sine curve is the set (see Figure 1c)

$$S = \{(x, \sin(1/x)) : x \in (0, 1]\} \cup \{(0, y) : y \in [-1, 1]\}$$

Proposition 3.2. The closed topologist's sine curve and the line segment cannot be distinguished by Σ_2^0 invariants.

Proof. Using Theorem 3.1, we simply have to show that S and the line segment strongly approximate each other. We work in the plane \mathbb{R}^2 , which embeds in Q. Let $I = \{(x, 0) : x \in [0, 1]\}$ and $\epsilon > 0$. We define continuous functions $f_n : I \to \mathbb{R}^2$ as follows (see Figure 1):

$$f_n(x,0) = \begin{cases} (x,\epsilon\sin(\frac{1}{x})) & \text{if } x \ge \frac{1}{n\pi}, \\ (x,0) & \text{if } x \le \frac{1}{n\pi}. \end{cases}$$

Each f_n is an ϵ -function, and the sets $f_n(I)$ converge to a copy of S, vertically scaled by ϵ . Therefore, I strongly approximates S.



Figure 1: ϵ -deformations of I converging to a copy of S

We now prove that $S \leq I$. Let $\epsilon > 0$. We build a continuous ϵ -function $f : S \to \mathbb{R}^2$. Let $k \in \mathbb{N}$ be large so that $a_k := \frac{1}{\pi/2 + 2k\pi} < \epsilon$. Note that $\sin(\frac{1}{a_k}) = 1$. Let

$$f(x,y) = \begin{cases} (0,y) & \text{if } 0 \le x \le a_k, \\ (x - a_k, y) & \text{if } x \ge a_k. \end{cases}$$

The function f satisfies $d(f, id) \leq a_k < \epsilon$ and f(S) is homeomorphic to I, so $S \leq I$. \Box

3.2 Cylinder

We show another application of Theorem 3.1. If X is a topological space, its **cylinder** is the space $X \times [0, 1]$ endowed with the product topology.

Theorem 3.2. If X is compact connected and not a singleton, then $X \leq X \times [0, 1]$.

Proof. The argument is illustrated in Figure 2.

We first build a sequence of continuous functions $f_n: X \to [0,1]$ such that for every nand every $x \in X$, $f_n(\overline{B}(x,\frac{1}{n})) = [0,1]$. As X is compact and not a singleton, there exists a finite set $F_n \subseteq X$ of cardinality at least 2 such that for every $x \in X$, there exists $y \in F_n$ with $d(x,y) < \frac{1}{2n}$. Let $\delta < \frac{1}{2n}$ be such that $d(y,y') \geq 2\delta$ for all distinct $y, y' \in F$.



Figure 2: Illustration of the proof of Theorem 3.2

Let us check that the function $f_n(x) = \min(d(x, F_n)/\delta, 1)$ satisfies the sought condition. For $x \in X$, let $y \in F_n$ be such that $d(x, y) < \frac{1}{2n}$. One has $\overline{B}(y, \delta) \subseteq \overline{B}(x, \frac{1}{n})$. For $z \in \overline{B}(y, \delta)$, $d(z, F_n) = d(z, y)$ takes values in $[0, \delta]$, and it must take all the possible values: if $v \in [0, \delta]$ is not reached, then X is covered by the two disjoint open sets B(y, v) and $X \setminus \overline{B}(y, v)$, which are non-empty because they contain y and some $y' \in F_n$ respectively, so X is disconnected, contradicting the assumption about X. Therefore, $f_n(\overline{B}(x, \frac{1}{n})) = [0, 1]$.

Now, we assume that X is already embedded in $(0,1)^{\mathbb{N}}$ and consider its copy $X_0 = \{0\} \times X$. For $\epsilon > 0$, we then define $g_n : X_0 \to Q$ as $g_n(0,x) = (\epsilon f_n(x), x)$. The sets $g_n(X_0)$ are ϵ -deformations of X_0 , converging to $[0,\epsilon] \times X$, which is a copy of $[0,1] \times X$. \Box

For instance, every Σ_2^0 invariant satisfied by the circle is also satisfied by the (geometric) cylinder, and every Σ_2^0 invariant satisfied by the line segment is satisfied by the filled square.

However, the cylinder and the filled square satisfy the Σ_2^0 invariant "having dimension at least 2" (Proposition 3.1), which is not satisfied by the circle and the line segment.

3.3 Wedge sum

We give another application of Theorem 3.1, showing that Σ_2^0 invariants cannot separate certain spaces, at least in one direction.

A pointed topological space is a pair (X, x_0) where X is a topological space and $x_0 \in X$ is a distinguished point. If (X, x_0) and (Y, y_0) are pointed spaces, then their wedge sum is the space $X \vee Y$ obtained by attaching X and Y at their distinguished points, which are identified.

Theorem 3.3. Let (X, x_0) and (Y, y_0) be pointed compact spaces. If Y is connected and x_0 is not isolated in X, then $X \leq X \vee Y$.

The idea of the proof is that for any $\epsilon > 0$, one can deform a small neighborhood of x_0 and make it converge to a small copy of Y of size ϵ .

Proof. Let $\overline{0} = (0, 0, 0, ...) \in Q$ be the sequence with only zeroes. First observe that when embedding a pointed space (Z, z_0) in Q, one can always send z_0 to $\overline{0}$ because Q is homogeneous, i.e. every point of Q can be sent to any other point by some homeomorphism $f: Q \to Q$ (Theorem 1.6.6 in [vM01]).

Let $Q_0 = \{x \in Q : \forall i, x_{2i} = 0\}$ and $Q_1 = \{x \in Q : \forall i, x_{2i+1} = 0\}$. These subspaces of Q are homeomorphic to Q: for i = 0 (resp. 1), the homeomorphism $f_i : Q \to Q_i$ just inserts zeroes at even (resp. odd) positions. By the previous observation, we can embed X in Q_0 and Y in Q_1 , so that $x_0 = y_0 = \overline{0}$.

As $x_0 = \overline{0}$ is not isolated in X, the function $d(x,\overline{0})$ takes infinitely many values in $\overline{B}(\overline{0}, 1/n) \cap X$ for any n. Therefore, the set

$$P_n := \{ nd(x,\overline{0}) : x \in \overline{B}(\overline{0}, 1/n) \cap X \}$$

is infinite. It implies by Proposition 2.4 that P_n approximates any connected space (either P_n contains a non-trivial connected component, or it contains infinitely many connected components, in which case P_n approximates every space). In other words, there exist continuous functions $h_n : P_n \to Q_1$ such that $h_n(P_n)$ converge to Y. Moreover, we can assume w.l.o.g. that if $1 \in P_n$, then $h_n(1) = \overline{0}$.

Let $\epsilon > 0$. We define $g_n : X \to Q$ as follows:

$$g_n(x) = \begin{cases} x & \text{if } 2/n \le d(x,\overline{0}), \\ (nd(x,\overline{0}) - 1)x & \text{if } 1/n \le d(x,\overline{0}) \le 2/n, \\ \epsilon h_n(nd(x,\overline{0})) & \text{if } d(x,\overline{0}) \le 1/n. \end{cases}$$

The function g_n is well-defined: if $d(x,\overline{0}) = 2/n$, then $g_n(x) = x$, if $d(x,\overline{0}) = 1/n$, then $g_n(x) = \overline{0}$. As the three pieces on which g_n is defined are closed and cover X, g_n is continuous (pasting lemma, [Mun00, Theorem 18.3]).

If n is sufficiently large, then one has $d(g_n(x), x) \leq 2\epsilon$ for all $x \in X$. When $n \to \infty$, $g_n(X)$ converges to $X \lor Y$.

4 Separating finite topological graphs

In order to understand which spaces can be separated by Σ_2^0 invariants, we focus on a restricted class of spaces, namely the finite topological graphs. We obtain a characterization of when a graph strongly approximates another graph. As a consequence, we show that Σ_2^0 invariants can separate any two graphs that are not homeomorphic. However, the result is not symmetric: for instance, for each $n \geq 2$ there is a Σ_2^0 invariant satisfied by the star with n + 1 branches and not by the star with n branches, but not vice-versa.

4.1 Modulus of local connectedness

Proving that a space X can be separated from a space Y by some Σ_2^0 invariant is not always easy, because only few classical invariants are Σ_2^0 , and we often need to design specific ones. Theorem 3.1 gives an alternative strategy, by showing that X does not strongly approximate Y. The proof of the theorem then provides a Σ_2^0 invariant (namely, having a copy which is a limit of ϵ -deformations of X), but this approach has two limitations:

- It may not be simple to show that X does not strongly approximate Y,
- The argument is not *effective*: it yields a $\sum_{n=1}^{\infty} 2^n$ invariant, but not a Σ_2^0 one.

In this section, we introduce a quantitative measure of local connectedness of the space, which can be used to separate many spaces and is effective, i.e. computable in some weak sense. We will see in the next section that it induces a complete invariant for the class of finite connected graphs.

If (X, d) is a metric space, $x \in X$ and r > 0, then we consider the open ball $B(x, r) = \{y \in X : d(x, y) < r\}$ and the closed ball $\overline{B}(x, r) = \{y \in X : d(x, y) \le r\}$.

Definition 4.1. Let (X, d) be a compact connected metric space. Its modulus of local connectedness is the function $\eta_X : (0, +\infty) \to [0, +\infty)$ defined by

 $\eta_X(r) = \min\{s \in [0, +\infty) : \forall x \in X, B(x, r) \text{ is contained in}\}$

a connected component of $\overline{B}(x,s)$.

The fact that it is indeed a min and not only an inf will be proved in the next proposition. A compact subset of Q inherits the metric from Q. Let $\mathcal{K}_{\text{conn}}(Q)$ be the space of connected compact subsets of Q. The next result tells us that the inequality $\eta_X(r) > s$ can be effectively tested.

Proposition 4.1. If X is compact connected, then the function η_X is well-defined and nondecreasing. Moreover, the set $\{(X, r, s) : \eta_X(r) > s\}$ is Σ_1^0 in the product space $\mathcal{K}_{\text{conn}}(Q) \times (0, +\infty) \times [0, +\infty)$.

Proof. We first show that $\eta_X(r)$ is well-defined, i.e. is indeed a min rather than an inf. Let $S_X(r)$ be the set involved in the definition of $\eta_X(r)$, namely

$$S_X(r) = \{s \in [0, +\infty) : \forall x \in X, B(x, r) \text{ is contained in } \}$$

a connected component of $\overline{B}(x,s)$.

This set is non-empty as it contains $d = \operatorname{diam}(X)$. Indeed, for all $x \in X$ one has $\overline{B}(x, d) = X$ which is connected by assumption. Note that if $s \in S_X(r)$ and s < s', then $s' \in S_X(r)$, so $S_X(r)$ is a half-line $[a, +\infty)$ or $(a, +\infty)$. We show that the second alternative is impossible. Assume that $S_X(r) = (a, +\infty)$ for some a. There exists $x \in X$, such that B(x, r) is not contained in a connected component of $\overline{B}(x, a)$. Equivalently, there exist two disjoint open sets U, V that both intersect B(x, r), and such that $\overline{B}(x, a) \subseteq U \cup V$. As X is compact, one still has $\overline{B}(x, a + \epsilon) \subseteq U \cup V$ for sufficiently small $\epsilon > 0$. Therefore, $a + \epsilon \notin S_X(r)$, which is a contradiction. As a result, $S_X(r) = [a, +\infty)$ for some a, and $\eta_X(r) = a = \min S_X(r)$.

If $s \in S_X(r)$ and r' < r, then $s \in S_X(r')$ because $B(x, r') \subseteq B(x, r)$, so $r \mapsto \eta_X(r)$ is non-decreasing.

We now show that $\{(X, r, s) : \eta_X(r) > s\}$ is Σ_1^0 . Let $(q_i)_{i \in \mathbb{N}}$ be a computable enumeration of the rational points of Q, and $(r_i)_{i \in \mathbb{N}}$ be a computable enumeration of the positive rational numbers. Let $(U_i)_{i \in \mathbb{N}}$ be a computable enumeration of the finite unions of balls centered at rational points, with rational radii.

Claim 4.1. One has $\eta_X(r) > s$ iff there exist $i, j, k, l, m, n \in \mathbb{N}$ such that $U_i \cap U_j = \emptyset$, X intersects $B(q_k, r_l), r_m + r_l < r, r_n > s + r_l$ and X intersects both $B(q_k, r_m) \cap U_i$ and $B(q_k, r_m) \cap U_j$ and $X \cap \overline{B}(q_k, r_n)$ is contained in $U_i \cup U_j$.

Proof. One has $\eta_X(r) > s$ if and only if there exist $x \in X$ and two disjoint open sets U, V such that $X \cap B(x, r)$ intersects both U and V and $X \cap \overline{B}(x, s)$ is contained in $U \cup V$.

If such x, U, V exist, then by compactness of $X \cap \overline{B}(x, s)$, one can assume that U and V are finite unions of rational balls U_i and U_j . The same conditions holds if one replaces r and s by sufficiently close rational numbers $r_m < r$ and $r_n > s$. Let then r_l be smaller than $r - r_m$ and $s - r_n$. The same conditions hold if one replaces x with a sufficiently close rational point q_k , and we can take $d(q_k, x) < r_l$ so that X intersects $B(q_k, r_l)$ at x.

Conversely, assume the existence of such i, j, k, l, m, n. Let $x \in X \cap B(q_k, r_l)$. One has $B(q_k, r_m) \subseteq B(x, r)$ and $\overline{B}(x, s) \subseteq \overline{B}(q_k, r_n)$ so B(x, r) intersects $X \cap U_i$ and $X \cap U_j$, and $X \cap \overline{B}(x, s) \subseteq U_i \cup U_j$, so $\eta_X(r) > s$.

As a result, $\{(X, r, s) : \eta_X(r) > s\}$ is the union, over all i, j, k, l, m, n such that $U_i \cap U_j = \emptyset$, of

$$\{X : X \text{ intersects } B(q_k, r_l) \text{ and } B(q_k, r_m) \cap U_i \text{ and } B(q_k, r_m) \cap U_j, \\ \text{and } X \cap \overline{B}(q_k, r_n) \subseteq U_i \cup U_j\} \times (r_m + r_l, +\infty) \times [0, r_n - r_l).$$

This is a Σ_1^0 set in the product topology (note that $X \cap \overline{B}(q_k, r_n) \subseteq U_i \cup U_j$ is equivalent to $X \subseteq U_i \cup U_j \cup \overline{B}(q_k, r_n)^c$, which is a Vietoris open set).

Observe that X is locally connected if and only if $\inf_{r>0} \eta_X(r) = 0$. It implies that being locally connected is Π_3^0 , as proved in [DS20].

The key property of the modulus of local connectedness is that it does not increase much under ϵ -deformations.

Proposition 4.2. If Y is a limit of ϵ -deformations of X, then

$$\eta_Y(r) \le \eta_X(r+2\epsilon) + \epsilon.$$

Proof. We first show the inequality when Y is an ϵ -deformation of X. Let $f: X \to Y$ be a surjective ϵ -function. Let r > 0 and $s < \eta_Y(r)$. There exist $x, y \in Y$ such that d(x, y) < r and x, y are not in the same connected component of $\overline{B}(x, s) \cap Y$. Let $x' = f^{-1}(x)$ and $y' = f^{-1}(y)$. One has $d(x', y') < r + 2\epsilon$, and x', y' are not in the same connected component of $\overline{B}(x', s-\epsilon) \cap X$. Indeed, if $C \subseteq \overline{B}(x', s-\epsilon) \cap X$ is connected and contains x', y',

then $f(C) \subseteq \overline{B}(x,s) \cap Y$ is connected and contains x, y, giving a contradiction. As a result, $\eta_X(r+2\epsilon) > s-\epsilon$. As this inequality holds for any $s < \eta_Y(r)$, we obtain the result.

Let $s = \eta_X(r + 2\epsilon) + \epsilon$. We have just shown that the set $\{Y : \eta_Y(r) \le s\}$ contains all the ϵ -deformations of X. It also contains their limits, because it is closed in the Vietoris topology by Proposition 4.1.

4.2 Effective vs non-effective classes

We now show that if X is a finite topological graph, then the Σ_2^0 invariants are no more expressive than the Σ_2^0 invariants for separating X from other spaces. We use the modulus of local connectedness to define a family of Σ_2^0 invariants, which are as powerful as arbitrary Σ_2^0 .

Theorem 4.1. Let X be a finite topological connected graph and Y a compact space. If there exists a Σ_2^0 invariant satisfied by X and not Y, then there exists a Σ_2^0 such invariant.

We will need the following lemma, which comes from the fact that one can embed any finite connected graph in Q in such a way that every ball is connected.

Lemma 4.1. Every finite topological connected graph has a copy X_0 in Q such that $\eta_{X_0}(r) \leq r$ for all r > 0.

Proof. Our goal is to embed a finite topological graph G in the Hilbert cube, in such a way that for every $x \in G$, every ball $B(x,r) \cap G$ is connected. It implies that $\eta_G(r) = r$. We recall that the Hilbert cube Q comes with the metric $d(x,y) = \sum_i 2^{-i} |x_i - y_i|$.

Let $V = \{0, ..., n\}$ be the vertex set. A vertex $i \in V$ is realized as

$$u_i = (0, \dots, 0, 2^{i-n-1}, 0, \dots) \in Q.$$

An edge (i, j) is realized as the line segment between u_i and u_j , i.e. the set of convex combinations of u_i and u_j . Therefore, any point of G can be expressed as $\lambda u_i + (1 - \lambda)u_j$ for some $i \neq j$ and $\lambda \in [0, 1]$.

We now describe the metric d inherited from Q. Let d_0 be the metric on any edge, identified with [0, 1] with the Euclidean metric.

Claim 4.2. • If x, y belong to the same edge, then $d(x, y) = 2^{-n} d_0(x, y)$.

- If x, y belong to adjacent edges with common vertex u, then one has $d(x, y) = 2^{-n} \max(d_0(x, u), d_0(y, u)).$
- If x, y belong to non-adjacent edges, then $d(x, y) = 2^{-n}$.

Proof. We write any $x \in Q$ as $x = (x_0, x_1, ...)$. Note that if $x \in G$ then $x_i = 2^{i-n-1}x'_i$ for some $x'_i \in [0, 1]$, and that $d(x, y) = \sum_i 2^{-i}2^{i-n-1}|x'_i - y'_i| = 2^{-n-1}\sum_i |x'_i - y'_i|$.

Let $x = \lambda u_i + (1 - \lambda)u_j$ and $y = \mu u_k + (1 - \mu)u_l$, with $i \neq j$ and $k \neq l$. There are three cases

- If k = i and l = j, then $d(x, y) = 2^{-n-1}(|\lambda \mu| + |\lambda \mu|) = 2^{-n}|\lambda \mu|$,
- If $k \neq i$ and l = j, then $d(x, y) = 2^{-n-1}(\lambda + \mu + |\lambda \mu|) = 2^{-n} \max(\lambda, \mu)$,
- If i, j, k, l are pairwise distinct, then $d(x, y) = 2^{-n-1}(\lambda + (1-\lambda) + \mu + (1-\mu)) = 2^{-n}$,

and the symmetric cases are similar. Note that $d_0(x, u_i) = \lambda$.

Let $x, y \in G$. We show that if $d(x, y) \leq s$, then there is a path from x to y in $\overline{B}(x, s)$. If $s \geq 2^{-n}$, then $\overline{B}(x, s) = G$ is connected. If $s < 2^{-n}$, then x and y belong either to the same edge or to two adjacent edges. In the first case, the segment from x to y is contained in $\overline{B}(x, s)$. In the second case, if u is the common vertex of the adjacent edges, then the segments from x to u and from u to y are contained in $\overline{B}(x, s)$. Therefore, $\overline{B}(x, s)$ is connected.

Proof of Theorem 4.1. If Y is disconnected, then "being connected" is a Π_1^0 invariant satisfied by X but not by Y. Now assume that Y is connected.

We assume that X is a simple graph (otherwise we subdivide it to make it simple). We fix a copy $X_0 \subseteq Q$ of X in Q provided by Lemma 4.1, so $\eta_{X_0}(r) \leq r$. What is important is not the exact upper bound on η_{X_0} , but the fact that it is computable. For rational $\alpha > 0$, we define an invariant \mathcal{P}_{α} as follows:

$$Y \in \mathcal{P}_{\alpha} \iff$$
 there exists $\beta < \alpha$ and a copy Y_0 of Y such
that $d_H(X_0, Y_0) < \beta$ and $\eta_{Y_0}(r) \le r + 3\beta$ for all $r > 0$.

Lemma 4.2. If α is rational, then \mathcal{P}_{α} is a Σ_2^0 invariant.

The idea of the proof is that we can replace arbitrary copies by a countable set of copies, so that the second-order existential quantifier becomes a first-order quantifier, like in the proof of Theorem 3.1, Claim 3.1.

Proof. There exists a computable sequence $(f_i)_{i \in \mathbb{N}}$ of injective functions $f_i : Q \to Q$, which is dense in the space of continuous functions from Q to Q (Lemma 4.13 in [AH23]). We show that the definition of \mathcal{P}_{α} would be equivalent if one replaces arbitrary copies Y_0 with copies $f_i(Y)$.

Assume the existence of β and Y_0 witnessing $Y \in \mathcal{P}_{\alpha}$. Let $f: Y \to Y_0$ be a homeomorphism, than we extend to a continuous function $f: Q \to Q$. Let $\epsilon < \alpha - \beta$, $\beta' = \beta + \epsilon < \alpha$ and f_i be ϵ -close to f. Note that $f_i \circ f^{-1}: Y_0 \to Q$ is an ϵ -function, so $f_i(Y) = f_i \circ f^{-1}(Y_0)$ is an ϵ -deformation of Y_0 . As a result,

$$\eta_{f_i(Y)}(r) \le \eta_{Y_0}(r+2\epsilon) + \epsilon \le r+3\beta',$$

which proves that we can restrict the copies of Y to copies of the form $f_i(Y)$. Therefore, the quantification over Y_0 can be replaced by a quantification over $i \in \mathbb{N}$. One can moreover

assume that β is rational. The rest of the formula defining \mathcal{P}_{α} is made of Σ_1^0 and Π_1^0 statements by Proposition 4.1 (note that X_0 is explicitly computable, so $d_H(X_0, f_i(Y))$ can be computed), so the whole expression is Σ_2^0 .

Note that $X \in \mathcal{P}_{\alpha}$ for every $\alpha > 0$. We show that for the purpose of separating X from other spaces, the family $(\mathcal{P}_{\alpha})_{\alpha \in \mathbb{Q}_{>0}}$ of Σ_2^0 invariants is as powerful as the entire family of all Σ_2^0 invariants.

We assume that $Y \in \mathcal{P}_{\alpha}$ for every $\alpha > 0$, and prove that X strongly approximates Y, implying that every Σ_2^0 invariant satisfied by X is satisfied by Y, by Theorem 3.1. Let $\epsilon > 0$. We show that some copy Y_0 of Y is a limit of ϵ -deformations of X_0 . The idea is that we can cut X_0 into small line segments, and that each line segment has ϵ -deformations converging to a small connected subset of Y_0 . Putting all the pieces together, X_0 has ϵ -deformations converging to Y_0 . Let us now give the details.

Let $\delta > 0$ be small (technically, we want $\delta < \epsilon/14$). We subdivide each edge of X by adding new vertices, so that the new edges have diameters $< \delta$. Let V and E be the new sets of vertices and edges respectively. For an edge e, let $U_e = \{x \in Q : d(x, e) < \delta\}$ be its open δ -neighborhood. One has diam $(U_e) < 3\delta$. Let Y_0 be a copy of Y witnessing that $Y \in \mathcal{P}_{\delta}$. As $d_H(X_0, Y_0) < \delta$, Y_0 is covered by the U_e 's and intersects each U_e .

One has $\eta_{Y_0}(3\delta) \leq 6\delta$. Each $Y_0 \cap U_e$ has diameter $\langle 3\delta$, so it is contained in a connected set $C_e \subseteq Y_0$ of diameter $\leq 12\delta$ by definition of η_{Y_0} (choose $x \in Y_0 \cap U_e$ and observe that $Y_0 \cap U_e \subseteq Y_0 \cap B(x, 3\delta)$ is contained in a connected component of $Y_0 \cap \overline{B}(x, 6\delta)$). As the U_e 's cover Y_0 , one has $Y_0 = \bigcup_e C_e$.

For each vertex $v \in V$, we choose a point $y_v \in Y_0$ such that $d(v, y_v) < \delta$. For each edge e that is incident to v, one has $y_v \in Y_0 \cap U_e \subseteq C_e$.

For each edge e, C_e is connected so it is approximated by the segment e, as in the proof of Proposition 2.4. In other words, there exist continuous functions $f_n^e : e \to Q$ such that $f_n^e(e)$ converge to C_e when n grows. Moreover, we can make sure that if v is an endpoint of e, then $f_n^e(v) = y_v$. Therefore, for each n, the functions f_n^e can be concatenated into one continuous function $f_n : X \to Q$, sending each v to y_v . The sets $f_n(X)$ converge to Y when n grows.

We finally show that for sufficiently large n, f_n is an ϵ -function. Let e be an edge and vone of its endpoints. One has diam $(e) < \delta$ and diam $(C_e) \le 12\delta$, so for $x \in e$ and $y \in C_e$,

$$d(x, y) \le d(x, v) + d(v, y_v) + d(y_v, y)$$

$$< \delta + \delta + 12\delta < \epsilon.$$

If n is sufficiently large, then $f_n^e(x)$ is close to C_e , so $d(x, f_n^e(x)) < \epsilon$.

The Σ_2^0 invariants defined in the proof are rather *ad hoc* and not very intuitive. It would be interesting to find a more natural family of Σ_2^0 invariants achieving the same effect.

4.3 Separating graphs

In this section, we prove that two non-homeomorphic finite topological graphs can be separated by some Σ_2^0 invariant. First, we show that a graph G can be separated from another graph H precisely when H cannot be contracted to G, even after subdivision.

In a graph, an edge contraction consists in removing an edge and identifying its two endpoints. A contraction is a sequence of edge contractions. A subdivision consists in adding new vertices inside the edges, and splitting the edges accordingly (it does not change the topology of the graph). We refer to $[HvHL^+14]$ for details on contractions.

We allow graphs with multiple edges and loops. Topologically, it makes no difference because subdividing such a graph results in simple graph which is topologically the same. However, edge contractions can create multiple edges and loops.

Theorem 4.2. Let G, H be finite connected graphs that are not singletons. The following statements are equivalent:

- There is a Σ_2^0 invariant satisfied by G but not H,
- No subdivision of H can be contracted to G.

Note that the same result holds for \sum_{2}^{0} invariants, by Theorem 4.1. For the proof, we use the following characterization of contractibility from [HvHL⁺14].

Lemma 4.3. G is a contraction of H iff there exists a surjective map $\varphi : V_H \to V_G$ such that:

- Each $\varphi^{-1}(s)$ is connected in G,
- For every pair of distinct vertices s, t of G, (s, t) is an edge of G iff there exists an edge e = (u, v) of H such that $\varphi(u) = s$ and $\varphi(v) = t$.

Proof of Theorem 4.2. The statement can be reformulated this way: if G, H are finite topological graphs, then one has $G \preceq H$ iff some subdivision of H can be contracted to G.

One implication is easy: if G is obtained from H by an edge contraction, then we show that G strongly approximated H.

Let e be the edge of H that is contracted and w be the vertex of G replacing e. Consider a small neighborhood of w, which is a star. For any $\epsilon > 0$, one can ϵ -deform this star, so that the endpoints are fixed and a small region around w is collapsed to an edge. The result of this ϵ -deformation is a graph in which w has been replaced by an edge; in other words, it is a copy of H. Therefore, G strongly approximates H.

We now prove the other direction, which is much more involved. Let us first discuss why subdividing H cannot be avoided. Consider the following example: H is just two vertices joined by an edge, and G is the subdivision of H obtained by adding a vertex in the middle. In that case, G and H are topologically the same, however H cannot be contracted to G,

because it has less vertices. The best we can hope is that some subdivision of H can be contracted to G.

The strategy is as follows:

- We fix a copy of G, a sufficiently small $\epsilon > 0$ and a copy of H which is a limit of ϵ -deformations of G,
- We subdivide H so that each edge of H is entirely contained in the ϵ -neighborhood of some edge of G,
- On each edge e of G, we choose a point a_e which is far from the vertices of both G and H,
- We show that a_e is close to exactly one edge of H,
- When removing the points a_e from G, one is left with a disjoint union of stars centered at the vertices of G,
- We show that H can be contracted to G, by sending each vertex v of H to the center of the star that is closest to v (and using Lemma 4.3).

We fix a copy of G given by Lemma 4.1. For simplicity of notation, we assume that the distances between two distinct vertices is 1 (rather than 2^{-k} , where k + 1 is the number of vertices of G; the argument could be easily adapted by adding a scaling factor). One has $\eta_G(r) \leq r$. Let n be the number of vertices of H.

Let e be an edge of G. There must be a point a in the middle third of e which is far from every vertex of H. More precisely, take n + 1 points regularly distributed on the middle third of e. They are all at distance at least $\frac{1}{3n}$ from each other. A vertex of H can be $\frac{1}{6n}$ -close to at most one of these points, so by the pigeonhole principle, some of these points is $\frac{1}{6n}$ -far from V_H . We choose such a point for each edge e of G and call it a_e .

Let $\delta \leq \frac{1}{6n}$ and $\epsilon < \delta/6$. Let G_n be $\epsilon/2$ -deformations of G converging to a copy of H. We call this copy H for simplicity.

We now subdivide H, so that every edge of H is entirely close to an edge of G.

Claim 4.3. We can subdivide H so that each edge of H is entirely contained in the ϵ -neighborhood of some edge of G.

Proof. The ϵ -neighborhoods of the edges of G form an open cover of the compact set H. Let μ be a Lebesgue number of the cover. We subdivide each edge of H so that the new edges have diameters smaller than μ , and are therefore contained in the ϵ -neighborhood of some edge of G.

From now on, we assume that H has been subdivided (note that n is still the original number of vertices of H, so δ and ϵ do not have to be redefined and there is no circularity).

Claim 4.4. For every $a \in G$, if $d(a, V_H) > \delta$ then $B(a, \epsilon)$ intersects H in exactly one edge of H.

Proof. First, $B(a, \epsilon)$ intersects H, because the Hausdorff distance between G and H is less than ϵ , and $a \in G$.

As *H* is a limit of ϵ -deformations of *G*, one has $\eta_H(r) \leq \eta_G(r+2\epsilon) + \epsilon = r + 3\epsilon$, so $\eta_H(2\epsilon) \leq 5\epsilon$.

If x and y are two points of $B(a, \epsilon) \cap H$, then $d(x, y) < 2\epsilon$ so x and y are connected in $\overline{B}(x, 5\epsilon) \cap H \subseteq B(a, \delta) \cap H$. This set is disjoint from V_H , so the only way for x and y to be connected in that set is that they belong to the same edge of H.

In particular, for each edge e of G, $B(a_e, \epsilon)$ intersects exactly one edge f_e of H. Conversely,

Claim 4.5. Each edge f of H can intersect at most one ball $B(a_e, \epsilon)$. In particular, the map $e \mapsto f_e$ is injective.

Proof. If $e' \neq e$, then $d(a_{e'}, e) \geq 1/3$, so $B(a_{e'}, \epsilon) \cap N_{\epsilon}(e) = \emptyset$ (assuming $\epsilon \leq 1/6$). Because of the preliminary subdivision of H, each edge f of H is contained in $N_{\epsilon}(e)$ for some edge e of G, therefore it does not intersect $B(a_{e'}, \epsilon)$ for $e' \neq e$, so it can only intersect $B(a_e, \epsilon)$.

It follows that the map $e \mapsto f_e$ is injective. Let $e' \neq e$. As $f_{e'}$ intersects $B(a_{e'}, \epsilon)$, it cannot intersect $B(a_e, \epsilon)$. As f_e intersects this ball, one has $f_{e'} \neq f_e$.

Let e be an edge of G. As e is convex, the function $r: Q \to e$ sending a point x to the closest point on e is continuous. Its restriction to f_e is a continuous function $h: f_e \to e$. The edge e is a copy of the unit interval [0, 1] and therefore inherits its natural ordering \leq (there are two possible orientations, we choose an arbitrary one).

Let u, v be the endpoints of f_e .

Claim 4.6. On the edge e, a is between h(u) and h(v).



Figure 3: An impossible case: x_b and x'_b are close, but not closely connected

Proof. We assume that $a < h(u) \le h(v)$ and derive a contradiction. The other cases where h(u) and h(v) are on the same side of a are symmetric, and the same argument applies.

As $f_e \subseteq N_{\epsilon/2}(e)$, one has $d(h(x), x) \leq \epsilon/2$ for all $x \in f_e$.

Let $x_a \in f_e$ be such that $d(x_a, a) \leq \epsilon/2$, implying that $d(h(x_a), a) \leq \epsilon$.

Let $b \in e$ be the middle point between a and h(u). One has $d(a, h(u)) \ge d(a, u) - d(u, h(u)) > \delta - \epsilon/2$ so $d(a, b), d(b, h(u)) > \delta/2 - \epsilon/4$.

As $a < b < h(u) \leq h(v)$, there exist x_b, x'_b such that $u < x_b < x_a < x'_b < v$ such that $h(x_b) = h(x'_b) = b$. One has $d(x_b, x'_b) \leq \epsilon$, so x_b and x'_b must be connected in $H \cap \overline{B}(x_b, 5\epsilon)$. In H, the only paths that connect x_b and x'_b cross x_a or u. Therefore, this ball must contain x_a or u. However, $d(x_b, x_a) \geq d(b, a) - \epsilon$ and $d(x_b, u) \geq d(b, h(u)) - \epsilon$. They are both larger than 5ϵ if ϵ is sufficiently small, $\epsilon < \delta/14$.

Let $A = \{a_e : e \in E_G\}$. The set $G \setminus A$ is a disjoint union of stars centered at the vertices of G. Each vertex v of H is ϵ -close to some point y_v of G, which belongs to one of these stars (indeed, $y_v \notin A$ because $d(v, A) > \delta > \epsilon$).

We use Lemma 4.3 and define a map $\varphi: V_H \to V_G$ sending $v \in H$ to the center of the star of y_v .

Let f = (u, v) be an edge of H such that $\varphi(u) \neq \varphi(v)$. f is ϵ -close to an edge e = (s, t)of G. e is 1/3-far from every star centered at a vertex other than s and t, so f is $(1/3 - \epsilon)$ far from these stars. As u and v are ϵ -close to the stars of $\varphi(u)$ and $\varphi(v)$ respectively, one must have $\varphi(u) = s$ and $\varphi(v) = t$ (or the symmetric case). As a result, $(\varphi(u), \varphi(v))$ is the arc e of G.

If e = (s, t) is an edge of G, then $f_e = (u, v)$ satisfies $\varphi(u) = s$ and $\varphi(v) = t$ because u, v stand on opposite sides of a_e .

Claim 4.7. Each $\varphi^{-1}(s)$ is connected.

Proof. Let $u, v \in \varphi^{-1}(s)$. y_u and y_v belong to the star S_s centered at s. The star is connected, so there exists a sequence of points $y_0, y_1, \ldots, y_k \in S_s$ with $y_0 = y_u, y_k =$ y_v and $d(y_i, y_{i+1}) < \epsilon$. As G and H are $\epsilon/2$ -close in the Hausdorff metric, there exist $x_0, \ldots, x_k \in H$ such that $d(x_i, y_i) \leq \epsilon/2$. One has $d(x_i, x_{i+1}) < 2\epsilon$, so x_i and x_{i+1} are connected in $\overline{B}(x_i, 5\epsilon) \cap H$. Therefore, u and v are connected in $N_{5\epsilon}(S_s) \cap H$, i.e. there is a path from u to v in H contained in $N_{5\epsilon}(S_s)$. All the vertices of this path are close to S_s , so they belong to $\varphi^{-1}(s)$. As a result, u and v are connected in $\varphi^{-1}(s)$.

Example 4.1. For instance, every Σ_2^0 invariant satisfied by the line segment is satisfied by the graphs that have a bridge, i.e. an edge whose removal disconnects the graph. Every Σ_2^0 invariant satisfied by the circle is satisfied by the graphs that have a cycle. Moreover, the line segment and the circle can be separated by some Σ_2^0 invariant, in both directions.

A consequence of this result is that non-homeomorphic graphs can be separated by some Σ_2^0 invariant, because they cannot be contracted to each other, even after subdivisions.

Theorem 4.3. If G, H are non-homeomorphic finite connected graphs, then there exists a Σ_2^0 invariant satisfied by one of them but not the other.

Proof. In order to apply Theorem 4.2, we need to show that G and H cannot be contracted to each other, even after subdivisions. We assume that some subdivision of G contracts to H and some subdivision of H contracts to G, and prove that G and H must be homeomorphic. It results from the following simple observations:

- In a graph G = (V, E), the sum $S = \sum_{v \in V: \deg(v) \ge 3} \deg(v)$ does not change when subdividing and does not increase when contracting an edge. Therefore, it must be the same in G, H and all the intermediate graphs of the contractions,
- If an edge contraction does not change S, then the resulting graph is homeomorphic to the original graph (the only exception would be if G is an edge or a loop, contracted to a point; but a point cannot be contracted back to G).

Let us prove these observations.

Let G be a graph, e = (u, v) be an edge, and let H be the new graph after contraction and w be the new vertex replacing u and v. Let S_G, S_H be the corresponding sums.

We first show that $S_H \leq S_G$. If e is a loop, i.e. u = v, then $\deg(w) = \deg(u) - 2$, so $S_H \leq S_G$. If e is not a loop, i.e. $u \neq v$, then $\deg(w) = \deg(u) + \deg(v) - 2$. If $\deg(u), \deg(v) \leq 2$, then $\deg(w) \leq 2$ and $S_H = S_G$; if $\deg(u) \geq 3$ and $\deg(v) \leq 2$ (or vice-versa), then $\deg(w) \leq \deg(u)$ so $S_H \leq S_G - \deg(u) + \deg(w) \leq S_G$; if $\deg(u), \deg(v) \geq 3$ then $\deg(w) < \deg(u) + \deg(v)$ so $S_H = S_G - \deg(u) - \deg(v) + \deg(w) < S_G$.

We now assume that $S_H = S_G$ and show that unless G is an edge or a loop, H is homeomorphic to G. First, e is not a loop: otherwise, as G is not a loop and is connected, one has $\deg(u) \ge 3$ so $S_H \le S_G - \deg(u) + \deg(w) \le S_G - 2$.

Claim 4.8. At least one of the vertices u and v has degree 2.

Proof. First, it is not possible that both $\deg(u)$ and $\deg(v)$ are at least 3, because in that case one would have $S_H = S_G - 2$. Therefore, swapping u and v if needed, one has $\deg(u) \leq 2$. If $\deg(u) = 1$, then $\deg(v) = 2$: if $\deg(v) = 1$ then G, which is connected, is just an edge; if $\deg(v) \geq 3$, then $S_H \leq S_G - 1$.

Assume that u has degree 2, the other case is symmetric. It means that u is adjacent to exactly one other edge f. The contraction of e simply shortens the line segment consisting of e and f, so H is homeomorphic to G.

Remark 4.1. It would be interesting to extend Theorem 4.3 to larger families of spaces. A natural such family is given by the finite simplicial complexes, which are higher-dimensional generalizations of graphs. However, the analog statement fails. For instance, if X is the disk and Y is the wedge of two disks, then these spaces are 2-dimensional simplicial complexes that are not homeomorphic, however they cannot be separated by Σ_2^0 invariants (it is

not difficult to show that they strongly approximate each other). It raises the following question: what is the level of complexity needed to separate 2-dimensional finite simplicial complexes? what about the *n*-dimensional ones?

5 Complexity of recognizing the line segment

Each particular space X induces a topological invariant, which is "being homeomorphic to X". The obvious descriptive complexity of this invariant is Σ_1^1 , because it is formulated using an existential quantifier over the continuous set of homeomorphisms.

However, a classical theorem from Descriptive Set Theory, due to Miller [Kec95, Theorem 15.14] and Ryll-Nardzewski [BK96, Theorem 2.3.4] implies that the complexity of recognizing a particular space is always Borel. For the interested reader, this result states that when a Polish group continuously acts on a Polish space, each orbit is Borel. In our context, the Polish group of homeomorphisms from Q to itself continuously acts on the Polish space of compact subsets of Q, and the set of copies of a space is an orbit of the group action, and is therefore Borel by this theorem.

This result opens up the following research program: for each concrete space X, what is the exact complexity of recognizing X? We answer this problem when X is the line segment.

Theorem 5.1. Being homeomorphic to [0,1] is Π_3^0 -complete.

The proof that it is Π_3^0 heavily relies on the following characterization of the line segment. If X is connected, then we say that a point $x \in X$ is a **noncut** point if $X \setminus \{x\}$ is still connected.

Theorem 5.2 ([Eng89, Theorem 6.3.8]). The line segment is the unique compact connected metrizable space which does not contain 3 noncut points.

The descriptive complexity of testing this property directly is a priori high, because it is quantified over the points of X. However, when X is locally connected, we show that it induces another characterization which is much simpler to test.

Definition 5.1. Let X be a compact metrizable space. We say that 3 points $x, y, z \in X$ are **independent** if they have pairwise disjoint neighborhoods U_x, U_y, U_z respectively, such that y, z are in the same connected component of $X \setminus U_x$, x, z are in the same connected component of $X \setminus U_x$.

For instance, in the circle, any 3 pairwise distinct points are independent. In the star with 3 branches, which is the space having the shape of the letter Y, the 3 endpoints of the branches are independent. However, the line segment has no set of 3 independent points.

As announced, this property is simple to test.

Proposition 5.1. Having at least 3 independent points is a Σ_2^0 invariant.

Proof. Let $\epsilon > 0$. Consider the property $\mathcal{P}_{\epsilon}(x, y, z)$: x, y are in the same connected component of $X \setminus B(z, \epsilon)$ and similarly for x, z and y, z. The property $\mathcal{P}_{\epsilon}(x, y, z)$ is Π_1^0 . Consider the property \mathcal{Q}_{ϵ} : $\exists x, y, z$ at distance $\geq \epsilon$ from each other, satisfying $\mathcal{P}_{\epsilon}(x, y, z)$. \mathcal{Q}_{ϵ} is obtained as an existential quantification over the compact set $\{(x, y, z) \text{ at distance } \geq \epsilon \text{ from each other} \}$, so \mathcal{Q}_{ϵ} is Π_1^0 as well [Pau16, Proposition 5.2]. The invariant is $\bigcup_{\epsilon > 0} \mathcal{Q}_{\epsilon}$ where we can take ϵ rational, so it is Σ_2^0 .

When X is locally connected, we show that if 3 points are noncut points, then they are independent.

Proposition 5.2. The line segment is the unique compact space that is connected, locally connected and has no set of 3 independent points.

Proof. Let X be compact, connected, locally connected and not homeomorphic to [0, 1]. We show that it has 3 independent points. By Theorem 5.2, X has at least 3 noncut points x, y, z. Using the assumption that X is locally connected, we show that these points are independent. We use another important result from topology, Mazurkiewicz-Moore-Menger's theorem [Eng89, Theorem 6.3.11], stating that a compact connected metrizable space is locally connected if and only if it is locally path-connected. Therefore, X is locally path-connected. As a result, $X \setminus \{z\}$ is also locally path-connected, so its connected components are path-connected [Mun00, Theorem 25.5]. As x and y belong to the same connected component of $X \setminus \{z\}$, they are connected by a path in $X \setminus \{z\}$. That path is a compact set that does not contain z, so its complement is a neighborhood U_z of z, and x, y are in the same connected component of $X \setminus \{u_z\}$. The same argument applied to x and to y shows that x, y, z are independent.

We can now prove the main result of this section.

Proof of Theorem 5.1. The invariant expressed by Proposition 5.2 is Π_3^0 . We show that recognizing the line segment is Π_3^0 -hard, by reducing the following Π_3^0 -complete problem: given an infinite binary matrix $M = (M_{i,j})_{i,j \in \mathbb{N}}$, decide whether all its rows contain only finitely many 1's. To each matrix M we associate a compact set $X_M \subseteq [0,1]^2$ which is homeomorphic to the line segment if and only if all the rows of M contain finitely many 1's. The function sending M to X_M is computable, i.e. there is a program that can compute approximations of X_M in the Hausdorff metric, given an access to M as oracle. The construction is illustrated in Figure 4.

First, to each row *i* we associate a compact set $X_i \subseteq [0,1]^2$ which is a copy of [0,1] if *s* contains finitely many 1's, and a copy of the topologist's sine curve *S* otherwise. It is defined as follows. Let $a_j = 2^{-j}$ and $m_j = \frac{a_j + a_{j+1}}{2}$. X_i contains the line segment from (0,0) to (0,1); for each *j*, if $M_{i,j} = 0$, then add a line segment from $(a_j,1)$ to $(a_{j+1},1)$; if $M_{i,j} = 1$, add a line segment from $(a_j, 1)$ to $(a_{j+1}, 1)$.

Note that when X_i is a copy of the line segment, its endpoints are (0,0) and (1,1).



Figure 4: Illustration of the Π_3^0 -completeness of the line segment

We then define the space X_M by concatenating the sets X_i , scaled and translated so that their endpoints match. Let $f_i : \mathbb{R}^2 \to \mathbb{R}^2$ be the function that scales by a factor 2^{-i-1} and then translates by the vector $(1 - 2^{-i}, 1 - 2^{-i})$. We define $X = \{(1, 1)\} \cup \bigcup_{i \in \mathbb{N}} f_i(X_i)$.

If every row of M contains finitely many 1's, then X_M is a concatenation of copies of the line segment, so it is itself a copy of the line segment. If some row contains infinitely many 1's, then X_M is not locally connected, so it is not a copy of the line segment. \Box

The proof of the result has another consequence: every locally connected space which is not homeomorphic to [0, 1] can be separated from [0, 1] by a Σ_2^0 invariant. Again, in this context Σ_2^0 invariants are no more expressive than Σ_2^0 invariant.

Corollary 5.1. If X is locally connected but not homeomorphic to [0,1], then there exists a Σ_2^0 invariant satisfied by X but not by [0,1].

Proof. Either X is disconnected or has at least 3 independent points. This property is a Σ_2^0 invariant which is not satisfied by [0, 1].

Note that the local connectedness assumption cannot be dropped. The closed topologist's sine curve S is not locally connected, and cannot be separated from [0,1] by $\sum_{i=1}^{0} 2^{i}$ invariants.

6 Future directions

The results presented in this article are the first steps of a much broader research program. We list a few important problems left for future investigations.

We have built *ad hoc* Σ_2^0 invariants separating finite topological graphs.

Problem 1. Is there a "natural" class of Σ_2^0 invariants separating finite topological graphs?

The family of finite graphs is very restricted. The finite simplicial complexes, which are generalizations of graphs to higher-dimensions, provide a rich family of spaces. Problem 2. For each $n \in \mathbb{N}$, what is the minimal level of complexity of topological invariants that can separate any pair of *n*-dimensional finite simplicial complexes? In the 1-dimensional case of graphs, we saw that it is Σ_2^0 .

We have identified the descriptive complexity of recognizing the line segment.

Problem 3. What is the descriptive complexity of recognizing the circle? the disk? The same question can be raised for any particular compact space.

References

- [AH22] Djamel Eddine Amir and Mathieu Hoyrup. Computability of finite simplicial complexes. In Mikolaj Bojanczyk, Emanuela Merelli, and David P. Woodruff, editors, 49th International Colloquium on Automata, Languages, and Programming, ICALP 2022, July 4-8, 2022, Paris, France, volume 229 of LIPIcs, pages 111:1–111:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.
- [AH23] Djamel Eddine Amir and Mathieu Hoyrup. Strong computable type. To appear in Computability, 2023.
- [Ama15] Manuel Amann. Computational complexity of topological invariants. Proceedings of the Edinburgh Mathematical Society, 58(1):27–32, 2015.
- [AS16] Michał Adamaszek and Juraj Stacho. Complexity of simplicial homology and independence complexes of chordal graphs. *Computational Geometry*, 57:8–18, 2016.
- [Bec92] Howard Becker. Descriptive set theoretic phenomena in analysis and topology. In Haim Judah, Winfried Just, and Hugh Woodin, editors, Set Theory of the Continuum, pages 1–25, New York, NY, 1992. Springer US.
- [BK96] Howard Becker and Alexander S. Kechris. The Descriptive Set Theory of Polish Group Actions. London Mathematical Society Lecture Note Series. Cambridge University Press, 1996.
- [BMS15] Benjamin A. Burton, Clément Maria, and Jonathan Spreer. Algorithms and complexity for turaev-viro invariants. In Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann, editors, Automata, Languages, and Programming, pages 281–293, Berlin, Heidelberg, 2015. Springer Berlin Heidelberg.
- [DM23] Rodney G. Downey and Alexander G. Melnikov. Computably compact metric spaces. *Bulletin of Symbolic Logic*, pages 1–59, 2023.

- [DS20] Gabriel Debs and Jean Saint Raymond. The descriptive complexity of connectedness in Polish spaces. *Fundamenta Mathematicae*, 249(3):261–286, 2020.
- [Eng89] Ryszard Engelking. *General topology. Rev. and compl. ed.*, volume 6. Berlin: Heldermann Verlag, rev. and compl. ed. edition, 1989.
- [HvHL⁺14] Pinar Heggernes, Pim van't Hof, Benjamin Lévêque, Daniel Lokshtanov, and Christophe Paul. Contracting graphs to paths and trees. *Algorithmica*, 68(1):109–132, 2014.
- [IK21] Zvonko Iljazović and Takayuki Kihara. Computability of subsets of metric spaces. In Vasco Brattka and Peter Hertling, editors, Handbook of Computability and Complexity in Analysis. Springer, 2021.
- [Kec95] Alexander S. Kechris. *Classical Descriptive Set Theory*. Springer, January 1995.
- [LMN23] Martino Lupini, Alexander Melnikov, and Andre Nies. Computable topological abelian groups. *Journal of Algebra*, 615:278–327, 2023.
- [Mun00] J.R. Munkres. *Topology*. Featured Titles for Topology. Prentice Hall, Incorporated, 2000.
- [Pau16] Arno Pauly. On the topological aspects of the theory of represented spaces. Computability, 5(2):159–180, 2016.
- [Vau74] Robert Vaught. Invariant sets in topology and logic. Fundamenta Mathematicae, 82(3):269–294, 1974.
- [vM01] J. van Mill. The Infinite-Dimensional Topology of Function Spaces. North-Holland Mathematical Library. Elsevier Science, 2001.