INF556 Topological Data Analysis
Final Exam — 3 hours

December 19, 2024

Important:

- The exercises are independent of one another.
- The text of the exam is written in English. Your answers can be written indifferently in French or in English.
- Please keep in mind that the quality of your answers (completeness of the arguments and clarity of their exposition) will be key for the grading.
- All printed documents are allowed. By contrast, computers, cellphones, tablets, pocket calculators, etc., are forbidden.

1 Monsters Inc.

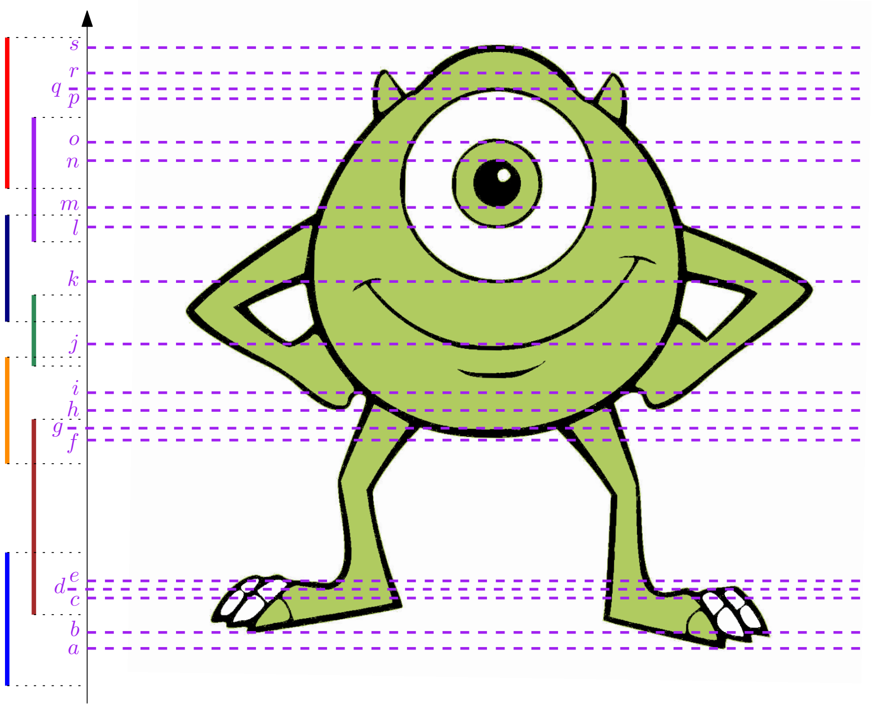

Figure 1: Meet Mike Wazowski (Bob Razowski in French).

Question 1. Draw the Reeb graph of the height function defined on the union of the green regions and their incident black curves in Figure 1.

Question 2. Draw the persistence diagram of the **relative part** of the extended persistent homology of the height function (for all homology degrees).

Question 3. Draw the Mapper of the height function, for the interval cover of the image of the function displayed in colors on the left of Figure 1.

2 Euler characteristic

Given a topological space X and a field $\mathbf{k}$, the *Euler characteristic* is the quantity:

$$\chi(X;\mathbf{k}) = \sum_{i=0}^{+\infty} (-1)^i \dim \mathrm{H}_i(X;\mathbf{k}).$$

Question 4. Show that χ is a topological invariant, that is: for any spaces X, Y that are homotopy equivalent, $\chi(X;\mathbf{k}) = \chi(Y;\mathbf{k})$.
Hint: look at what happens to each individual homology group.

Now we want to prove the Euler-Poincaré theorem:

Theorem 1. *For any simplicial complex X and any field $\mathbf{k}$:*

$$\chi(X;\mathbf{k}) = \sum_{i=0}^{+\infty} (-1)^i n_i(X),$$

where $n_i(X)$ denotes the number of simplices of X of dimension i.

For this we will use topological persistence. Consider an arbitrary filtration of X:

$$\emptyset = X_0 \subseteq X_1 \subseteq \cdots \subseteq X_m = X.$$

Assume without loss of generality that a single simplex σ_j is inserted at each step j:

$$\forall j = 1, \cdots, m, \ X_j \setminus X_{j-1} = \{\sigma_j\}.$$

Note that m is then equal to the number of simplices of X, that is:

$$m = \sum_{i=0}^{+\infty} n_i(X).$$

Let us apply the persistence algorithm to this simplicial filtration. Recall from the course that we have the following property:

Lemma 2. *At each step j, the insertion of simplex σ_j either creates an independent d_j-dimensional cycle (i.e. increases the dimension of $\mathrm{H}_{d_j}(X_{j-1};\mathbf{k})$ by 1) or kills a $(d_j - 1)$-dimensional cycle (i.e. decreases the dimension of $\mathrm{H}_{d_j-1}(X_{j-1};\mathbf{k})$ by 1), where d_j is the dimension of σ_j.*

Question 5. Using Lemma 2, prove Theorem 1.
Hint: you may proceed by induction on j, although this is not the only way to prove the result.

Question 6. Deduce that the Euler characteristic of a triangulable space is independent of the choice of field $\mathbf{k}$.

3 The Dunce Hat

Recall that the Dunce Hat is obtained by indentifying the three edges of a triangle as shown in Figure 2.

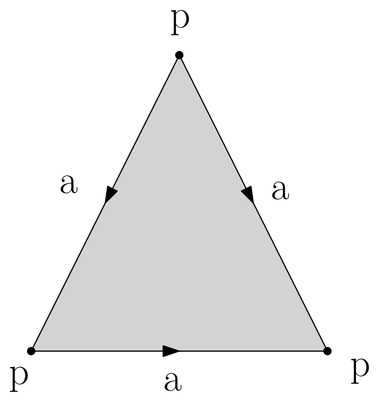

Figure 2: The Dunce Hat.

Question 7. Build a triangulation of the Dunce Hat (you may draw a picture to represent it). Beware that your triangulation must be a simplicial complex, not a general cell complex.

Question 8. Use your simplicial complex to compute the homology of the Dunce Hat over the field $\mathbb{Z}/2\mathbb{Z}$.
Hint: to avoid tedious calculations, you can proceed as in exercise 2: pick a filtration of your complex then apply the persistence algorithm; for each simplex σ_j inserted, use Lemma 2 to predict its effect on the homology (identify the created d_j-cycle or the killed $(d_j - 1)$-cycle).

4 Mayer-Vietoris sequences and the Jordan curve theorem

It is a well-known fact that any simple closed curve in the plane splits the plane into two connected components. The same holds for simple closed curves on the 2-sphere $\mathbb{S}^2$. This is the result we want to prove, modulo some mild regularity condition on the curve. For this we will use Mayer-Vietoris sequences: given two open sets $A, B \subseteq \mathbb{S}^2$, there exist morphisms in homology that make the following sequence **exact** (i.e. such that the image of each map is equal to the kernel of the next map in the sequence):

$$H_1(A \cap B) \xrightarrow{i_1} H_1(A) \oplus H_1(B) \xrightarrow{j_1} H_1(A \cup B) \xrightarrow{\partial_1} H_0(A \cap B) \xrightarrow{i_0} H_0(A) \oplus H_0(B) \xrightarrow{j_0} H_0(A \cup B) \xrightarrow{\partial_0} 0$$

This sequence is called the *Mayer-Vietoris long exact sequence* of A, B. The precise definition of the morphisms i_*, j_*, ∂_* is unimportant for the exercise. The notation $H_r(X)$ stands for the degree-r singular homology group of space X over some fixed field of coefficients (omitted in the notations). Note that the exactness of the sequence implies in particular that the morphism j_0 is surjective. Indeed, exact sequences of the following forms:

$$\cdots \longrightarrow 0 \longrightarrow X \xrightarrow{f} Y \longrightarrow \cdots$$

$$\cdots \longrightarrow X \xrightarrow{f} Y \longrightarrow 0 \longrightarrow \cdots$$

$$\cdots \longrightarrow 0 \longrightarrow X \xrightarrow{f} Y \longrightarrow 0 \longrightarrow \cdots$$

imply respectively that the morphism f is injective ($\operatorname{Ker} f = \operatorname{Im} 0 = 0$), surjective ($\operatorname{Im} f = \operatorname{Ker} 0 = Y$), and bijective.

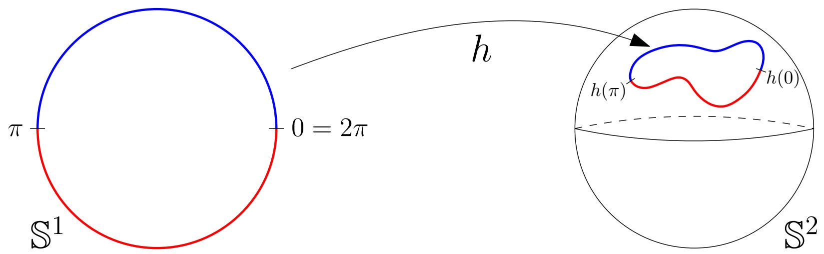

Figure 3: Embedding h of the circle $\mathbb{S}^1$ into the sphere $\mathbb{S}^2$. The circle is split into 2 arcs, namely: $[0, \pi]$ (in blue) and $[\pi, 2\pi]$ (in red).

Question 9. Given two open sets $A, B \subseteq \mathbb{S}^2$ such that $A \cap B = \emptyset$, use their Mayer-Vietoris sequence to prove that $H_*(A \cup B) \simeq H_*(A) \oplus H_*(B)$ for each $* \in \{0, 1\}$.

Question 10. Given two open sets $A, B \subseteq \mathbb{S}^2$ that are both homotopy equivalent to the disk $D^2 = \{x \in \mathbb{R}^2 : \|x\|_2 \leq 1\}$, prove that

$$\dim H_0(A \cap B) + \dim H_0(A \cup B) = 2 + \dim H_1(A \cup B).$$

Deduce the relation between the number of holes in $A \cup B$ and the number of path-connected components in $A \cap B$.

We will now prove the Jordan curve theorem. Let $h : \mathbb{S}^1 \to \mathbb{S}^2$ be a homeomorphism onto its image. We split the curve $h(\mathbb{S}^1)$ into 2 (non-disjoint) arcs, namely $h([0, \pi])$ and $h([\pi, 2\pi])$, as illustrated in Figure 3. To simplify the analysis we assume the following regularity condition on the curve:

(H) For every closed interval $I \subsetneq \mathbb{S}^1$, the space $\mathbb{S}^2 \setminus h(I)$ is homotopy equivalent to the punctured sphere $\mathbb{S}^2 \setminus \{\mathrm{pt}\}$ (where pt denotes an arbitrary point of $\mathbb{S}^2$).

This condition holds for instance when h is a diffeomorphism, i.e. when the curve is smooth.

Question 11. Show that $\dim H_0(\mathbb{S}^2 \setminus h(\mathbb{S}^1)) = 2$.

Note that the conclusion of the theorem still holds without hypothesis **(H)**. However, the proof contains more technicalities to handle the potentially erratic local behavior of the curve.

5 Decomposition of persistence modules

Let us fix a **finite** index set $T \subseteq \mathbb{R}$ of cardinality $\#T < +\infty$, and an arbitrary field of coefficients, and let us consider **pointwise finite-dimensional** persistence modules over T, i.e. persistence modules $\mathbb{V}$ such that $\dim V_t < +\infty$ for all $t \in T$.

A *submodule* $\mathbb{W}$ of a persistence module $\mathbb{V}$ is composed of subspaces $W_t \subseteq V_t$ for all $t \in T$, and of maps $w_t^{t'} = v_t^{t'}|_{W_t}$ for all $t \leq t' \in T$. In particular, $v_t^{t'}(W_t) \subseteq W_{t'}$ for all $t \leq t' \in T$. A simple example of submodule is the *null module* $\mathbb{W} = 0$ (defined by $W_t = 0$ for all $t \in T$ and $w_t^{t'} = 0$ for all $t \leq t' \in T$), which is a submodule of any module $\mathbb{V}$ over T.

Direct sums in the category are defined pointwise, that is: for any persistence modules $\mathbb{V}$ and $\mathbb{W}$ over T, the direct sum $\mathbb{V} \oplus \mathbb{W}$ is composed of the spaces $V_t \oplus W_t$ for all $t \in T$, and of the maps $v_t^{t'} \oplus w_t^{t'}$ for all $t \leq t' \in T$. We say that $\mathbb{V}$ and $\mathbb{W}$ are *summands* of the direct sum $\mathbb{V} \oplus \mathbb{W}$. In particular, we have $\mathbb{V} = 0 \oplus \mathbb{V} = \mathbb{V} \oplus 0$ for any persistence module $\mathbb{V}$, so $\mathbb{V}$ is always a summand of itself. A persistence module $\mathbb{V}$ is called *indecomposable* if its only summands are itself or the null module. The decomposition theorem that we saw in class asserts that the only indecomposable persistence modules are the so-called *interval modules* over T, and that (under some conditions on T or on the dimensions of the spaces) any module decomposes as a direct sum of interval modules.

Question 12. In the case where $\#T = 1$, show that every submodule is a summand. Deduce the decomposition theorem in the case $\#T = 1$.

In the general case, the result is more complicated to prove as submodules may not always be summands:

Question 13. In the case where $\#T \geq 2$, exhibit a counterexample showing that a submodule of a persistence module $\mathbb{V}$ over T may not always be a summand of $\mathbb{V}$.

The case $\#T = 2$ is somewhat simpler to handle though, as only one linear map is involved.

Question 14. Prove the decomposition theorem in the case where $\#T = 2$.

The general (finite) case involves more technicalities. First, you must show that your module actually decomposes into indecomposables:

Question 15. Show that every (pointwise finite-dimensional) module over T (finite) decomposes into indecomposable modules. You may proceed for instance by induction on the *total dimension* of $\mathbb{V}$, defined as the quantity $\sum_{t \in T} \dim V_t$.

Then, the structure of each indecomposable must be studied:

Question 16 (difficult). Show that every indecomposable (pointwise finite-dimensional) module over T (finite) is an interval module.

The decomposition theorem follows. Note that we have not considered the uniqueness of the decomposition here.