ARCHIVE OF MECHANICAL ENGINEERING

DOI: 10.24425/ame.2026.157610 2026, Vol. 73, No. 1, pp. 195–215

Janusz PRZEWOCKI [2,1], Maria MARCHWICKA, Anna WĄSIK [3,1], Bartosz REICHEL [4], Wiktor SIEKLICKI [5], Katarzyna GOCH [6,1], Krzysztof ŁUKASZUK [7,1]

Detection of daily activities from ear-worn accelerometer signals: a pilot study with topological data analysis

Received 24 September 2025, **Revised** 29 December 2025, **Accepted** 29 January 2026, **Published online** 12 March 2026

Keywords: accelerometer, human activity recognition, topological data analysis, persistent homology, machine learning

The purpose of this research was to develop and investigate a method for real-time lifestyle assessment, with the long-term goal of associating detected behaviours with menstrual cycle phases, eating habits and hydration. This was achieved by designing a system for detecting health-related activities – eating, drinking and smoking – using a custom ear-worn device equipped with an accelerometer. To account for confounding behaviours, speaking and a generic "other activities" category were also included. Fifteen prototype devices were built and tested in the laboratory, and nine were worn by healthy volunteers for data collection.

Raw accelerometer signals were segmented into overlapping windows and processed using topological data analysis (TDA) to extract shape-aware features. Time-delay embeddings were applied to transform signals into point clouds, from which persistence

✉ Wiktor Sieklicki, email: wiktor.sieklicki@pg.edu.pl

[1]Lifeconcept sp. z o.o., Olsztyn, Poland

[2]Institute of Mathematics, Faculty of Mathematics, Physics and Informatics, University of Gdańsk, Gdańsk, Poland

[3]Doctoral School at the Faculty of Mathematics, Physics and Informatics, University of Gdańsk, Gdańsk, Poland

[4]Institute of Applied Physics and Informatics, Faculty of Applied Physics and Mathematics, Gdańsk University of Technology, Gdańsk, Poland

[5]Institute of Mechanics and Machine Design, Faculty of Mechanical Engineering and Ship Technology, Gdańsk University of Technology, Gdańsk, Poland

[6]Departament of Social Medicine, Medical University of Gdańsk, Gdańsk, Poland

[7]Department of Obstetrics and Gynecology Nursing, Medical University of Gdańsk, Gdańsk, Poland

diagrams were computed using Vietoris-Rips and lower-star filtrations. Statistical descriptors derived from these representations included entropy measures, lifetime distributions, and topological complexity metrics.
On a dataset of 47 658 labelled windows collected from participants, a random forest classifier achieved 89% balanced accuracy on a held-out test split. Results demonstrate that TDA enables effective discrimination of mandibular and head movements captured by an ear-worn sensor. The source code and feature definitions are publicly available to support reproducibility.

1. Introduction

Human activity recognition (HAR) seeks to infer behaviours from sensor data. Most solutions such as smartwatches, smartbands, smartrings track activities connected to movement (physical activity), body signals (i.e., BBT, HRV) and their combinations (indirect assessment of sleep patterns, stress). These activities are common and clinically relevant for lifestyle and disease risk.

Previous work has linked dietary patterns [1, 2] and hydration [3, 4] with morbidity and cognitive function, while smoking remains a leading preventable cause of disease. Poor diet patterns contribute to approximately 11 million deaths worldwide annually [5] and are associated with obesity, type 2 diabetes, and cardiovascular disease [6–8]. Although the high energy intake is the most important risk factor, also eating frequency and speed may be relevant variables to be followed and considered [1, 2]. Research demonstrates that dietary self-monitoring significantly improves long-term weight management [9]. This suggests that automated detection of eating frequency, speed, and timing could provide beneficial behavioural feedback. Inadequate hydration impairs cognitive performance even at mild dehydration levels (1–2% body water loss) [10] and increases risks of kidney stones, dental problems [11, 12], and mortality [13, 14]. Both the amount of water consumed and drinking frequency are considered important factors influencing hydration [3, 4]. Smoking remains a leading preventable cause of disease, and objective feedback systems – such as carbon monoxide biomarker monitoring – have demonstrated efficacy in promoting cessation [15]. Research has shown that self-reported smoking behaviours and cigarette counts are prone to recall bias, particularly in individuals attempting to reduce or quit smoking [16].

The purpose of this research was to demonstrate that ear-worn accelerometer can effectively capture the subtle mandibular and head movements that characterise eating, drinking and smoking. Moreover, persistent homology features enable reliable discrimination between these activities.

In addition to eating, drinking and smoking, speaking also generates mandibular movements and produces acceleration patterns that partially overlap with those of the target activities. Therefore, speaking was included as a separate activity category in the classification, alongside a generic ”other activity” label.

The proposed system uses a device for activity recognition to provide detailed information on health-related behaviours and ultimately support timely interventions to promote healthier lifestyles.

In this framework, we use persistent homology, a topological data analysis (TDA) technique, as the method for classifying the ear-worn sensor signals. This approach characterises the global structure of time-series by capturing patterns that persist independently of noisiness and choice of temporal scale [17, chapter 59]. The topological features extracted via persistent homology offer complementary information about structure of the signal that is difficult to obtain with traditional time- or frequency-domain features [18]. Furthermore, these features are inherently robust to noise, as persistent homology emphasises enduring patterns in the data rather than transient fluctuations. In summary, our contributions are:

- a compact ear-mounted prototype for continuous temporomandibular joint (TMJ) movements monitoring,
- a labelled dataset comprising >47k windows from nine participants,
- a TDA-based pipeline that transforms short segments into persistence-derived features for multiclass classification.

This study is a part of a broader research and development initiative focused on exploring the potential of wearable technology for health-related applications. The overarching goal of the project was to develop a prototype wearable system capable of continuously measuring physical activity in the temporomandibular joint, along with other physiological parameters, to support fertility assessment and family planning. It may also be used as a hydration assessment tool, particularly in older adults. The system integrates hardware for signal acquisition and original software employing machine learning to analyse the collected data.

2. Persistent homology and its graphical representations

Topology is a branch of mathematics concerned with the study of shapes and, possibly, multidimensional, spatial structures. Topological data analysis is a framework that applies these ideas to datasets, aiming to uncover geometric structure that persist across scales and are robust to noise. Its conceptual foundation was laid by Carlsson in his seminal paper Topology and Data [19], which discussed a framework for using topological features – such as connected components, loops, and voids – to extract meaningful information from complex datasets. However, some of the core ideas underlying TDA, particularly persistent homology, were developed earlier in the context of computational topology [20].

Two key methodologies in TDA are the Mapper algorithm (and its multiple variants) and persistent homology. Mapper provides a simplified graphical summary of a high-dimensional dataset and has found applications in biomedicine, where it has been used to identify subgroups of patients and disease phenotypes [21]. In contrast, persistent homology focuses on computing topological invariants across multiple scales to detect stable features such as loops or voids in

data. These features are summarised in persistence diagrams or barcodes and offer a noise-resistant, coordinate-free representation of the underlying data structure (cf. Section 2.2).

Persistent homology is particularly well suited for time-series analysis [22, 23]. A common approach is to embed time-series data into higher-dimensional space using sliding windows [22] – forming point clouds whose geometry captures dynamics of the signal. Persistent homology can then identify and quantify recurrent behaviours, such as periodicity, by measuring persistence of loops in the embedded space [24]. This approach has proven effective in revealing signal structure that may not be evident through traditional methods.

In this paper, we adopt persistent homology to classify accelerometer signals captured by our ear-worn device. Topological features derived from sliding window embeddings serve as input to our machine learning models. This approach provides a principled way to extract shape-based signal features that are invariant to small deformations and robust to noise.

2.1. Homology theory

Homology is a classical and important concept in mathematics that dates back to the last decade of the 19th century [25]. In its persistent variant, it is one of the fundamental concepts of TDA. Here, we present its basic notions using simplicial complexes as the basis for its definition (this version is thus called *simplicial homology*).

2.1.1. Simplices and simplicial complexes

A simplex is a generalisation of geometric objects such as triangles and tetrahedra. They have several advantages over other types of "building blocks", e.g., have a clear geometric interpretation and work well with algebraic methods.

Definition 1 *A k-simplex is a convex hull spanned by $k+1$ affinely independent points $v_i \in \mathbb{R}^m$ and is denoted by:*

$$\sigma = [v_0, v_1, ..., v_k]. \tag{1}$$

Definition 2 *A simplicial complex K is a collection of simplices satisfying two conditions:*

1. *if a simplex σ is in K, then all faces of σ are also in K;*
2. *the intersection of two simplices in K is either empty or a simplex contained in K.*

2.1.2. Simplicial homology

To construct a simplicial homology groups, we need to endow a simplex with orientation. The orientation of the k-simplex is given by ordering of the vertices

in the simplex $[v_1, v_2, ..., v_{k+1}]$. Two orderings can define the same orientation if and only if they differ by an even permutation, thus there are only two possible orientations of simplex. Each oriented simplex is equal to the negative of the simplex with opposite orientation. Mathematically, it is stated as $[v_1, v_2, v_3] = -[v_2, v_1, v_3]$.

In simplicial homology, our aim is to identify cycles and determine whether a given cycle bounds a collection of higher-dimensional simplices: if it does not, we may intuitively think of the cycle as representing a "hole". To formalise the detection of such cycles in a simplicial complex, we use the concept of a simplicial k-chain, defined as follows.

Definition 3 *A k-chain is a finite weighted sum defined on all k-simplices within a complex K:*

$$\sum_{i=1}^{N} c_i \sigma_i, \tag{2}$$

where $c_i \in \mathbb{Z}$ and N is the number of k-dimensional simplices. The set of k-chains on K is denoted by $C_k(K)$. We use the short-hand notation $C_k(K) = C_k$ for a given K.

The boundary operator is a linear map that takes k-chains of a simplicial complex to their corresponding $(k-1)$-dimensional boundaries.

Definition 4 *Let $k > 0$ and $\sigma = [v_0, v_1, ..., v_k]$ be oriented k-simplex, viewed as a basis element of C_k. The boundary operator $\partial_k : C_k \to C_{k-1}$ is defined by:*

$$\partial_k(\sigma) = \sum_{i=0}^{k} (-1)^i [v_0, v_1, \ldots, v_{i-1}, v_{i+1}, \ldots v_k]. \tag{3}$$

For $k = 0$, we define $\partial_k = 0$ by convention.

Because simplices forms a basis of C_k, the boundary operator ∂_k can be defined on simplices and extended to all k-chains, giving the linear map $\partial_k : C_k \to C_{k-1}$.

Theorem 1 *For any k-chain $c_k \in C_k$ we have*

$$\partial_{k-1}(\partial_k(c_k)) = 0. \tag{4}$$

Let us focus on the most relevant information within the set of k-chains – two geometrically motivated subgroups: cycles and boundaries. The quotient of these subgroups provides the formal definition of a homology group.

Definition 5 *The k-dimensional cycles are given by:*

$$Z_k = \ker \partial_k, \tag{5}$$

where $\ker \partial_k$ is the kernel of the operator ∂_k.

Definition 6 *The boundaries are given by:*

$$B_k = im\ \partial_{k+1}, \tag{6}$$

where $im\ \partial_{k+1}$ *is the image of the operator* ∂_{k+1}.

Since composition of boundary operators is zero (see Theorem 1), we have the inclusion $B_k \subseteq Z_k$, which is fundamental property of the boundary group. We also note that groups C_k, Z_k and B_k – up to isomorphism – do not depend on orientation.

Definition 7 *The kth-homology group* H_k *is given by the quotient group:*

$$H_k = Z_k / B_k. \tag{7}$$

In simple terms, the group H_k is nonzero precisely when there exists k-cycles in the complex that are not boundaries. Intuitively, H_k captures the distinct k-dimensional "holes" in the complex.

2.2. Graphical summaries

Topological data analysis adapts homology theory for application to real-world datasets. However, most datasets consist of discrete point clouds that inherently lack topological properties such as connectedness or cycles. Therefore, the idea is to consider each point along with its neighbourhood of a growing radius, and monitor the properties of the union of such neighbourhoods rather than the points themselves (see Fig. 1). For computational purposes, the neighbourhoods are then replaced by their simplex counterparts defined below.

Definition 8 *Let* A *be a finite set in a metric space* (X, d). *The Vietoris-Rips complex of the set* A *for a fixed* $\varepsilon > 0$ *is the simplicial complex with vertices in* A *such that the set* $\{a_1, \ldots, a_n\} \subseteq A$ *is a simplex if* $d(a_i, a_j) < \varepsilon$ *for every* $1 \leqslant i, j \leqslant n$. *We will denote it by* $\mathcal{R}_\varepsilon$.

As shown in Fig. 1, the points in the dataset are surrounded by their respective balls, and then the centres of such balls are connected by a simplex if and only if the distance is sufficiently small. Following this construction we obtain the Vietoris-Rips complex. Obviously, the appearance of simplices is immediately related to the size of ε: the greater ε becomes, the less strict the criterion is, and the larger number of simplices appear. Persistent homology describes how the homology of a space changes as the scale parameter ε varies. Formally, it is obtained by applying the homology functor from Definition 7 to the Vietoris-Rips filtration defined below.

Definition 9 *Let* (X, d) *be a metric space. The Vietoris-Rips filtration is a sequence of Vietoris-Rips complexes* $(\mathcal{R}_\varepsilon)_{\varepsilon>0}$ *such that*

$$\mathcal{R}_{\varepsilon_0} \subseteq \mathcal{R}_{\varepsilon_1} \subseteq \ldots \subseteq \mathcal{R}_{\varepsilon_m}, \tag{8}$$

where $\varepsilon_0 < \varepsilon_1 < \ldots < \varepsilon_m$.

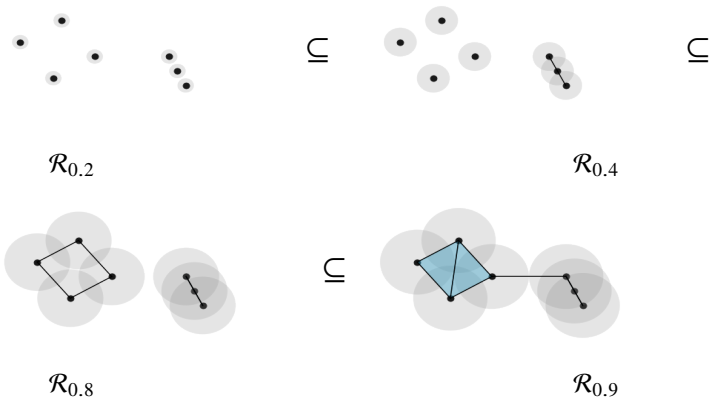

Fig. 1. Example of filtration

It is useful to interpret the growing parameter ε as *time*: as time passes, the complexes in the filtration change. More precisely, certain elements of homology groups appear and disappear or, as is often said in TDA, are born and die.

The Vietoris-Rips filtration is a main mathematical tool that we use in this paper to classify accelerometer data (cf. Section 3.3). From this filtration, one typically produces topological summaries that represent the births and deaths of homology classes associated with it. The most commonly used summaries are persistence diagrams, persistence landscapes [26] and barcodes [27]. Here, we focus on persistence diagrams, which we present in the following subsection

2.2.1. Persistence diagrams

Persistence diagrams (see Fig. 2) are topological summaries which we rely on the most in this paper. They depict births b and deaths d of homology classes presented in the form of points (b, d) on the plane. Since $b < d$ for every b, d, all points of the diagram are stored above the diagonal $d = b$. Formally, a persistence diagram is a multiset, i.e., it allows multiple occurrences of points (b, d), meaning that certain homology classes are born at the same time b and die at the same time d.

2.2.2. Lower-star filtration

Persistence diagrams generated from the lower-star filtration is another approach to summarise topological properties. This method is applied to real-valued functions rather than data clouds (it can also be extended to univariate time series via interpolation, which is the case for the accelerometer data considered here).

Given a function $f\colon X \to \mathbb{R}$ define the sublevel set of f as

$$A_\varepsilon(f) = \{x \in X\colon f(x) \leqslant \varepsilon\}. \tag{9}$$

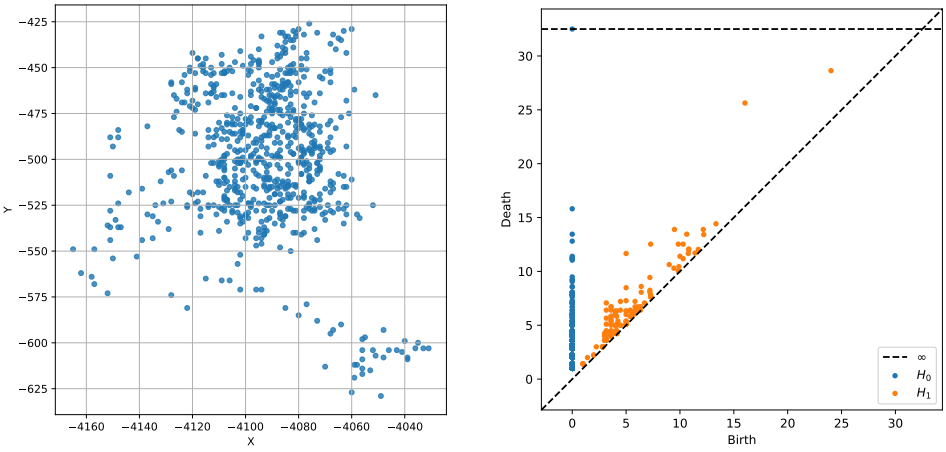

Fig. 2. An example of a point cloud generated from acceleration components of "drinking" event (shown on the left) and the corresponding persistence diagram (on the right)

It is essentially a preimage of the set $(-\infty, \varepsilon]$. Obviously, as ε increases, so do the sublevel sets. Therefore it is a filtration in a similar sense as the Vietoris-Rips counterpart: if $\varepsilon_1 < \varepsilon_2$ then $A_{\varepsilon_1}(f) \subseteq A_{\varepsilon_2}(f)$.

Intuitively, ε can be seen as the *sea level* and the filtration encodes what part of the domain of the function f is below the sea level (see Fig. 3). Therefore, the

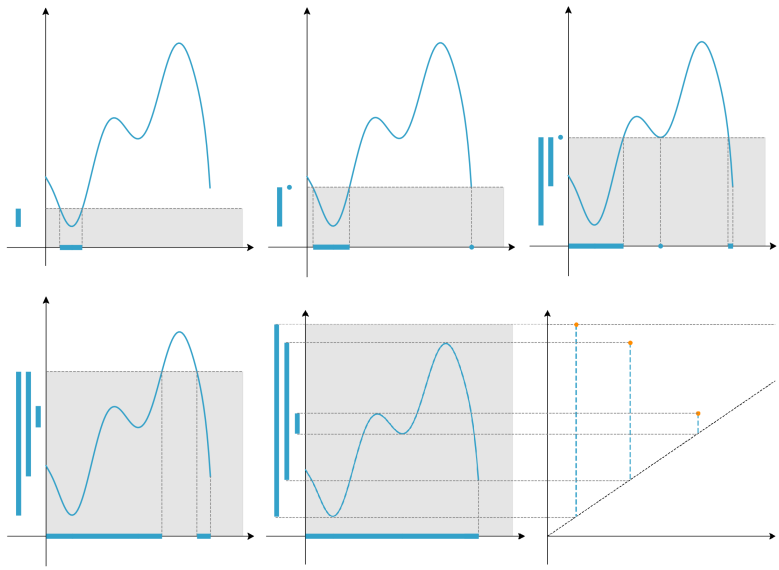

Fig. 3. Lower-star filtration of a polynomial function and its corresponding persistence diagram. The subplots illustrate the evolution of sublevel sets (grey regions) as the function values increase, highlighting the birth and death of connected components

nature of the filtration depends on the domain of the function. If the domain is $\mathbb{R}$, then each $A_\varepsilon(f)$ is a subset of $\mathbb{R}$ and hence the only homology group possible to examine is H_0. If f happens to be a multivariable function, then the filtration can encode cycles of a higher dimension.

3. Materials and methods

3.1. Ear wearable device with accelerometer

A compact, head-mounted prototype (Fig. 4a and b) has been developed to capture jaw and head motion, positioned superior to the ear and adjacent to the temporomandibular joint (TMJ). The core of this device is a custom printed circuit board (PCB) built around an ST32WB55 microcontroller (Fig. 4, d). This dual-core microcontroller is optimised for real-time, low layer operation and power management. The second PCB houses a 6-axis ISM330 Inertial Measurement Unit (Fig. 4c) , from which a 3-axis accelerometer is used to sample data at 416 Hz mode.

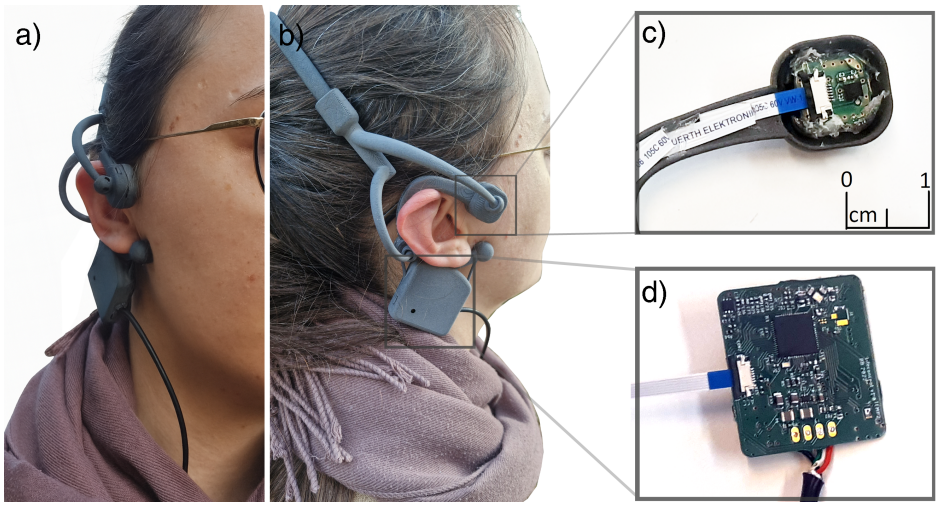

Fig. 4. Device positioned on the head (a) and (b), circuit board with ISM330 accelerometer mounted into the structure of the device (c), and circuit board with ST32WB55 microcontroller (d)

The custom board is connected via a USB cable to a Banana Pi M2 Zero single-board computer (SBC). The SBC runs multiple services to manage the data flow. The information flow within the experimental setup is presented in Fig. 5. In the current prototype configuration, the target device (comprising PCB.IMU and PCB.MCU) mounted on the participant's head is connected via a wired interface to a data acquisition unit (SBC). The participant interacts with buttons located on the controller (SBC.Controller) to annotate the activity being performed. A "data recorder" (SBC.Recorder) service gathers the time-stamped accelerometer data and

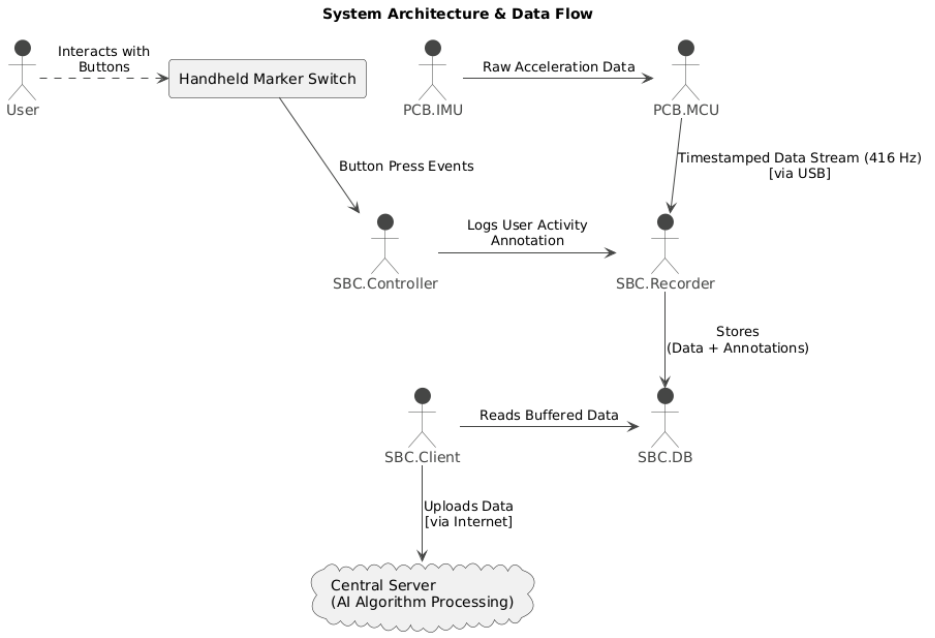

Fig. 5. System Architecture & Data Flow

stores it in a local database. This interaction is required during the data collection and model training phase.

Importantly, in the final system configuration, the SBC module will be omitted, and the head-mounted device will communicate wirelessly with a smartphone, eliminating the need for user interaction, meaning no buttons and no annotation required from the user. In the current version of the device, data stored locally at the SBC unit are transmitted to a cloud-based infrastructure, where the machine learning model is trained. Concurrently, a client service continuously checks for an internet connection and, when available, transmits the stored data to a central server for processing by AI algorithms. The entire system is powered by a portable power bank.

For user activity annotation, the system includes a handheld controller (*Marker Switch/SBC.Controller*) connected to the main module. The controller includes a set of buttons specifically designed to be operated by hand, allowing the user to easily select activities, even without having to look at the device. A dedicated "controller service" on the SBC monitors this input to log user activities annotated with the Marker Switch.

The housing of the prototype consists of two main parts created using different 3D printing technologies. The enclosure for the PCB, sensor and wiring is manufactured using Selective Laser Sintering (SLS) with Polyamide (PA12). This material is known for its high strength, stiffness, and suitability for functional parts. The SLS also provided the possibility to manufacture complicated shapes all-in-one

and without supports. The adjustable headband is produced using Fused Deposition Modeling (FDM) with a PET filament, a material selected for its flexibility and durability, while being easy to manufacture and inexpensive.

3.2. Data collection

To collect activity data for the classification of eating, drinking, smoking, and speaking, we conducted a user study involving fifteen participants equipped with a custom ear-worn wearable device. The device incorporated a triaxial accelerometer for motion capture and a set of physical buttons for real-time activity annotation.

Participants wore the device on the left ear, with the accelerometer positioned over the temporomandibular joint and stabilised using a headband. Four physical buttons were used to annotate the target activities, and one button was used to signal annotation errors. Participants manually labelled each activity using the corresponding button, according to the following standardized procedure:

- **Speaking**: Pressed at the start of speech and released during pauses longer than a few seconds.
- **Eating**: Pressed when food was placed in the mouth, held during chewing, and released after swallowing the last bite.
- **Drinking**: Pressed at the beginning of swallowing, held during the drinking episode, and released after the final sip.
- **Smoking**: Pressed at the onset of inhalation and released after exhalation.

In case of annotation errors (e.g., incorrect or multiple buttons pressed), participants were instructed to release all buttons and press the error button for two seconds.

Each participant underwent a brief training session to become familiar with the device and annotation protocol. Participants were allowed to engage in their normal activities (including walking, sitting, bending, etc.) during recordings and to carefully annotate their activities. This setup was specifically designed for data collection and model training and was largely driven by the time and budget constraints of the study. In subsequent development stages, the physical buttons are planned to be replaced by in-app confirmation functions within a mobile application. In the final product, model predictions will be verified only intermittently and randomly in a subset of users to support ongoing model refinement and validation under real-life conditions.

Nine women participated in the data acquisition stage of the study, each providing two recordings of approximately 15 minutes in duration. The most frequently observed class was “other activity” (45%), which included all behaviours other than the four target activities: eating, speaking, drinking, and smoking. Among the labelled activities, eating accounted for 21.7% of the total recording time, followed by speaking (16%), drinking (15%) and smoking (2.3%). The relatively low frequency of smoking is explained by the fact that only two out of nine participants were smokers.

Raw accelerometer signals do not provide an unambiguous representation of the activity performed by the user. To demonstrate this, representative time-series are presented for four activities, each shown along the x, y, and z axes (Fig. 6). Despite the richness of the recorded data, visual inspection does not allow reliable discrimination between activities. Therefore, the proposed methodology employs topological data analysis to extract informative structural features.

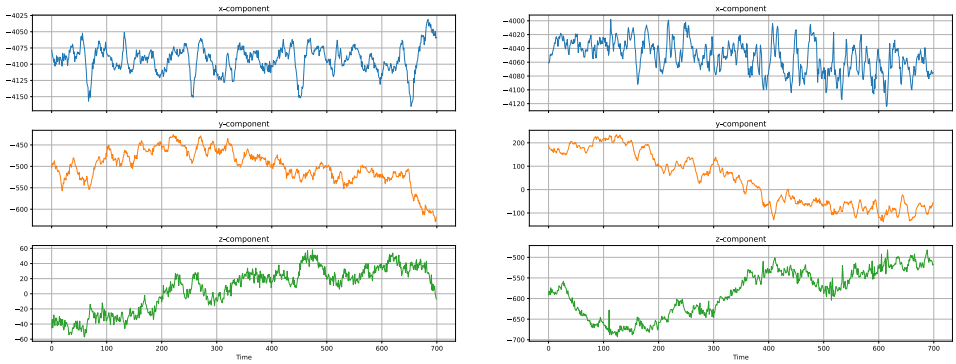

(a) Event: drinking

(b) Event: smoking

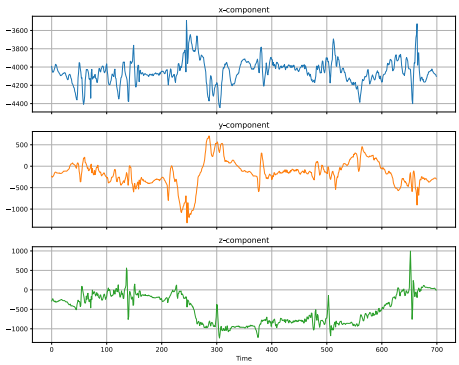

(c) Event: eating

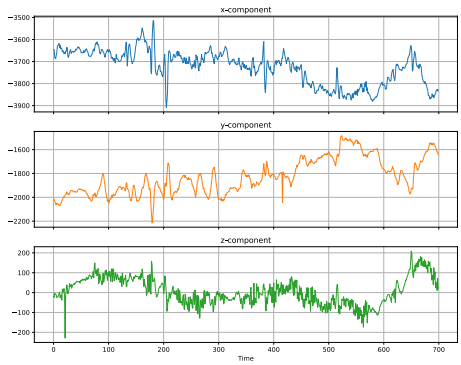

(d) Event: speaking

Fig. 6. Representative raw accelerometer time-series recorded along the x, y, and z axes for four activities. Activity labels were obtained from user annotations via button presses during data acquisition

3.3. Classification of signals from accelerometer using topological data analysis

Preparation of the input data for classification involved segmenting raw accelerometer signals into overlapping windows of 700 samples each, with a step size of 100 samples. This windowing strategy was designed to capture sufficient temporal context for activity recognition (about 2.3 seconds) while generating a

dense set of training examples. During the preprocessing, we excluded windows labelled as "error", as well as any windows that contained transitions between two or more activity classes, retaining only those segments where a single, unambiguous activity label could be assigned. This ensured that the classifier would learn from clearly defined examples of each activity type.

To classify time windows of accelerometer data corresponding to specific human activities, we used features extracted from persistence diagrams computed from time-delay embeddings. Each time window contained three univariate signals corresponding to the individual components of acceleration vector. For each window, we constructed three two-dimensional point clouds by pairing signal components: a_{xy}, a_{xz} and a_{yz}. These two-dimensional point clouds captured joint temporal dynamics across signal pairs. In some cases, we also found it useful to transform these two-dimensional signals back into univariate form by computing the Euclidean norm or extracting the polar angle (i.e., converting Cartesian coordinates to polar coordinates and projecting onto a single dimension).

Subsequently, each point cloud was lifted to a higher-dimensional space using a sliding window embedding described below, allowing for the extraction of topological features from the resulting spatiotemporal trajectories.

Let $x = \{x_1, x_2, \ldots, x_T\}$ denote a segment of the signal, where each $x_i \in \mathbb{R}^k$ represents a multivariate observation at time i. For example, $k = 2$ for original two-dimensional point clouds formed from signal pairs (e.g., a_{xy}), or $k = 1$ for univariate signals derived via transformations such as the Euclidean norm or polar angle.

Given the embedding dimension m and lag τ, we constructed a higher-dimensional point cloud P using a sliding window embedding:

$$P = \left\{ \left(x_t, x_{t+\tau}, x_{t+2\tau}, \ldots, x_{t+(m-1)\tau}\right) \mid t = 1, \ldots, T - (m-1)\tau \right\}. \tag{10}$$

Each element of $P \in \mathbb{R}^{k \cdot m}$. This embedding transforms the temporal structure of the signal into a spatial object, enabling the analysis of its topological properties such as connected components and loops.

In our analysis, we generated multiple types of point clouds from the segmented accelerometer signals to capture distinct geometric and temporal characteristics. The following constructions were used:

- **Two-dimensional point clouds** constructed directly from pairs of acceleration components: a_{xy}, a_{xz}, and a_{yz}, for each window.
- **Delay-embedded point clouds** created from the two-dimensional signals above using sliding window embeddings. For embedding dimension $m \in \{2, 5\}$ and $\tau \in \{0, 30\}$ the resulting point clouds lie in $\mathbb{R}^{2m}$, i.e., $\mathbb{R}^4$ or $\mathbb{R}^{10}$, respectively.
- **Three-dimensional point clouds** formed by appending time as an additional coordinate to the original 2D point clouds, yielding embeddings in $\mathbb{R}^3$. This construction preserves temporal ordering while capturing geometric structure.

- **Polar-transformed delay embeddings** derived from univariate signals obtained by transforming each 2D signal into scalar time series. Two transformations were used: the vector norm (magnitude) and the polar angle. Delay embeddings were then computed from these univariate signals using embedding dimensions $m \in \{2, 10\}$.

Persistence diagrams were computed from the point clouds described above using Vietoris-Rips (see Definition 9) filtrations in homology dimensions 0 and 1, implemented via the `ripser` v0.6.10 library [28]. Additionally, for signals transformed into polar coordinates, we applied a star filtration (cf. Section 2.2.2) to capture complementary topological information. In total, 22 persistence diagrams were generated for each time window, covering multiple signal transformations, homological dimensions, and filtration strategies.

From each persistence diagram, we extracted a set of 337 features describing the topological structure of the time window. These features included:

Descriptive Statistics. For birth times b_i, death times d_i, and lifetimes $l_i = d_i - b_i$, we computed:

- Mean: $\bar{l} = \frac{1}{n} \sum_i l_i$
- Variance: $\mathrm{Var}_l = \frac{1}{n} \sum_i (l_i - \bar{l})^2$

Counting Features. The number of homological features (i.e., the cardinality of each diagram) was recorded in each dimension.

Wasserstein Amplitude. A measure of the total topological persistence of a diagram. It is defined the Wasserstein distance [29] from the persistence diagram containing the diagonal only:

$$W_p(D) = \left(\sum_i |d_i - b_i|^p \right)^{1/p}. \tag{11}$$

In this work we used W_2 as a topological feature of a diagram.

Persistence Entropy. To measure the complexity of the persistence diagram D in an information-theoretic sense [30], we computed persistence entropy:

$$H(D) = -\sum_i \frac{l_i}{L} \cdot \log\left(\frac{l_i}{L}\right), \quad \text{where} \quad L = \sum_i l_i. \tag{12}$$

Vasicek's differential entropy estimator. We applied a classical non-parametric method for estimating the differential entropy of a continuous distribution, known as Vasicek's estimator [31]. This feature was computed for the sequence of birth times b_i or death times d_i extracted from the persistence diagram. Given a random sample $x = (x_1, x_2, \ldots, x_n)$, the estimator is defined as:

$$H_{\text{Vasicek}}(x) = \frac{1}{n} \sum_{i=1}^{n} \log \left(\frac{n}{2m} \left[x_{(i+m)} - x_{(i-m)} \right] \right), \tag{13}$$

where $x_{(i)}$ denotes the i-th order statistic of the sample and m is a window parameter controlling local neighborhood size. In our implementation, we set $m = 1$. This estimator provides an approximation of the differential entropy without requiring explicit knowledge of the underlying probability density function.

The dataset obtained after segmentation and cleaning contained the following number of time windows summarised in Table 1. The topological features were computed for each window and used as input to machine learning models for activity classification.

Table 1. Number of 700-sample windows per activity

Activity	Other activity	Eating	Speaking	Drinking	Smoking
Windows	21448	7634	10341	7133	1102

4. Results

The classification of temporomandibular joint activities was performed using a random forest (RF) classifier implemented in Python v3.10.12 with the `scikit-learn` v1.6.1 library [32]. Hyperparameter optimisation was conducted using the `Optuna` v4.3.0 framework [33], aiming to maximise the *balanced accuracy* metric. This metric was chosen due to the class imbalance in the dataset, particularly the under-representation of the smoking class. Balanced accuracy was computed via 5-fold cross-validation, where the training data were randomly split into five equally sized folds.

The dataset was divided into separate training (80% windows) and test subsets (20% windows); all reported performance metrics were computed on the test set to ensure unbiased evaluation. A total of 47658 time windows, each of 700 samples with an overlap of 600 samples, were analysed. Each window was represented by a set of topological features derived from persistence diagrams, as described in Section 3. These features captured statistical properties of birth times, death times, and lifetimes of homology classes, as well as Wasserstein amplitudes and entropy measures, across multiple dimensions and filtration strategies.

After feature extraction, Recursive Feature Elimination with Cross-Validation (RFECV) was applied to rank the importance of features. In our analysis, we

employed a version of RFECV implemented in the Python library `scikit-learn`. This method requires the classifier to compute feature importances, a criterion met by the random forest classifier. The algorithm operated as follows:

1. **Data Partitioning:** The data were divided into folds, where each fold used a subset of the data as test data and the remaining samples for training. The number of folds equalled the number of predefined partitions, ensuring that each subset was used exactly once for testing.
2. **Feature Elimination:** For each fold, the RFE algorithm began by iteratively removing features. The classifier was first fitted to compute feature importances, after which the least important feature was removed. The model's score was then calculated using the fold's test data. This process was repeated until only one feature remained.
3. **Score Averaging:** The scores calculated for each fold and each number of features during step 2 were averaged to obtain mean scores as a function of the number of features. The optimal number of features, n_{features}, was defined as the number with the highest mean score.
4. **Final Model Fitting:** The classifier was fitted over the entire dataset, and the n_{features} with the highest importance were selected.
5. **Iteration:** Steps 1–4 were repeated a predefined number of times or until only one feature remained.

This procedure also allowed us to assign a *feature score* to each selected feature. At each iteration in step 4, a subset of features was selected; whenever a feature appeared in this subset, its score was incremented by one. Consequently, features that persisted longer across iterations accumulated higher scores.

The most informative feature was the number of homology classes obtained via star filtration applied to the time series of the coordinate angle of the acceleration vector in the xz plane (see Fig. 7). We also observed an overrepresentation of entropy-related features, particularly those based on Vasicek's estimator of differential entropy computed with respect to the deaths of homology classes.

The final model achieved balanced accuracy 89.9% on the test set; more detailed metrics computed for each activity are presented in Table 2. The confusion matrix revealed high classification performance across all activities, with particularly strong discrimination for eating (recall 93.8%) and for all other activities with recall above 85% (see Fig. 8).

The most problematic activity was smoking, although this class also had the lowest number of windows in the dataset. The lowest precision (88%) was observed for the "other activity" class, indicating that the algorithm occasionally detected an activity when none was present. Most of these misclassifications occurred when other activity was recognised as drinking. Additional real-time tests suggested that, in some participants, swallowing saliva was misclassified as drinking. This highlights the need for post-processing steps to improve real-time recognition accuracy.

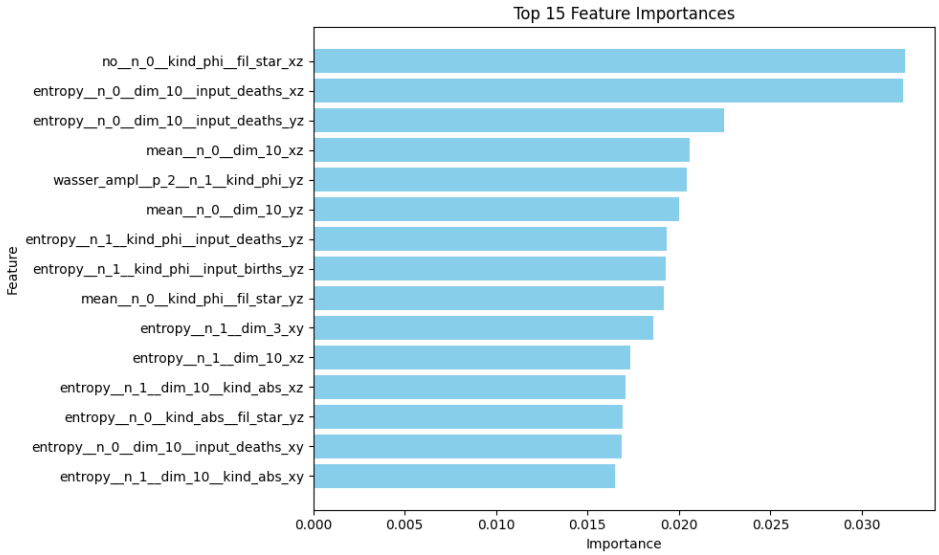

Fig. 7. Feature importances (Gini importances) computed from the final random forest model

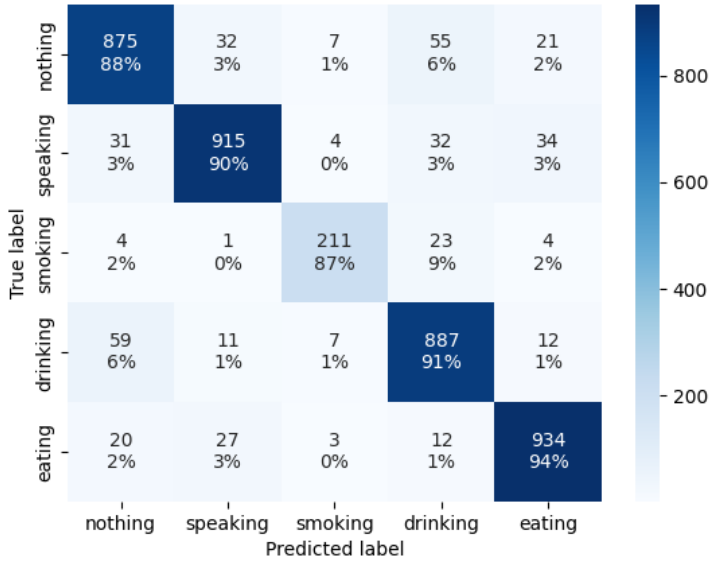

Fig. 8. Confusion matrix for the random forest classifier evaluated on the test set

These results demonstrate the feasibility of combining persistent homology-derived features with a mechanically simple, ear-worn accelerometer device for accurate recognition of temporomandibular joint-related activities.

Table 2. Classification metrics for each activity on the test set. Precision, recall, F1-score, and Matthews correlation coefficient (MCC) are reported per class. Overall macro-averaged values are shown in the last row

Activity|	Precision	Recall	F1-score	MCC
Eating	0.9294	0.9378	0.9335	0.9129
Drinking	0.8791	0.9088	0.8937	0.8612
Speaking	0.9280	0.9006	0.9141	0.8876
Smoking	0.9095	0.8683	0.8884	0.8820
Other activity	0.8847	0.8838	0.8843	0.8489
Overall (macro)	0.9061	0.8999	0.9028	0.8785

5. Conclusions

In this study, we demonstrated that topological data analysis provides a powerful framework for characterising subtle head and mandibular movements captured with an ear-worn accelerometer. By transforming raw signals into point clouds and computing persistence diagrams, we derived shape-aware descriptors that capture dynamic patterns beyond the reach of conventional feature extraction. Our classifier achieved a balanced accuracy of 0.89 across four activity classes, showing that TDA-derived features enable robust recognition of health-related behaviours: eating, drinking and smoking. Speaking was also detected for differentiation from other activities. These results highlight the potential of topology-driven approaches to enrich wearable sensing pipelines with interpretable and mathematically grounded representations.

Several limitations of this study should be acknowledged. First, the dataset comprised only nine participants (all female), which restricts the generalisability of the findings and may not capture inter-individual variability in head motion patterns, cultural eating habits, or sensor placement differences. Moreover, class imbalance was present, with smoking under-represented in the dataset. The evaluation was performed at the window level, which is appropriate given the small number of participants; however, in larger cohorts, subject-level validation strategies such as leave-one-group-out would provide a more rigorous assessment of external validity.

This work represents a preliminary proof-of-concept study. The results provide valuable insights into the feasibility of ear-worn sensing for lifestyle monitoring and demonstrate a strong potential of this approach for future health-tracking applications. Developing a fully deployable solution will require extensive validation in larger and more diverse populations, as well as substantial refinements to the device (miniaturization, improved user comfort, more attractive industrial design and/or extended functionality) to increase the likelihood of sustained or frequent use. The proposed technology could also be integrated into existing consumer products, such as bone-conduction headphones.

Research ethics statement. The study protocol was reviewed and approved by the Independent Bioethics Committee for Scientific Research at the Medical University of Gdańsk (Approval No. NKBBN/31/2023 issued on February 15, 2023). All procedures were conducted in accordance with the Declaration of Helsinki and relevant institutional guidelines. Prior to enrolment, each (adult) participant received written information about the study and provided written informed consent. Participation was voluntary, with the right to withdraw at any time without penalty. All data were pseudonymised prior to analysis to safeguard confidentiality and processed in compliance with applicable data protection regulations.

Funding. Co-funded by European Regional Development Fund under the Regional Operational Programme Intelligent Development 2014-2020. Grant No. POIR.01.01.01-00-0656_21

Conflicts of interest. Some of the authors are employed by Lifeconcept Ltd., including Janusz Przewocki, Anna Wasik, Bartosz Reichel, Wiktor Sieklicki, Katarzyna Goch and Krzysztof Lukaszuk. The device described in this work was developed within the framework of a research project conducted at Lifeconcept Ltd. The authors declare that these affiliations did not influence the objectivity of the study.

Code availability. The code used for preprocessing, feature extraction, and classification in this study is available at: https://github.com/jprzew/tmj_har_using_tda.

References

[1] T.E. Garcidueñas-Fimbres, I. Paz-Graniel, S.K. Nishi, J. Salas-Salvadó, and N. Babio. Eating speed, eating frequency, and their relationships with diet quality, adiposity, and metabolic syndrome, or its components. *Nutrients*, 13(5):1687, 2021. doi: 10.3390/nu13051687.

[2] G. Longo-Silva, M. de Oliveira Lima, A.K.P. Pedrosa, R. Serenini, P. de Menezes Marinho, and R.C.E. de Menezes. Association of largest meal timing and eating frequency with body mass index and obesity. *Clinical Nutrition ESPEN*, 60:179–186, 2024. doi: 10.1016/j.clnesp.2024.01.022.

[3] H. He, J. Zhang, N. Zhang, S. Du, S. Liu, and G. Ma. Effects of the amount and frequency of fluid intake on cognitive performance and mood among young adults in baoding, hebei, china: A randomized controlled trial. *International Journal of Environmental Research and Public Health*, 17(23):8813, 2020. doi: 10.3390/ijerph17238813.

[4] J. Frąckiewicz and K. Szewczyk. Is there an association between hydration status, beverage consumption frequency, blood pressure, anthropometric characteristics, and urinary biomarkers in adults? *Nutrients*, 17(6):952, 2025. doi: 10.3390/nu17060952.

[5] A. Afshin, P.J. Sur, K.A. Fay, L. Cornaby, G. Ferrara, J.S. Salama, E.C. Mullany, K.H. Abate, C. Abbafati, Z. Abebe, et al. Health effects of dietary risks in 195 countries, 1990–2017: a systematic analysis for the global burden of disease study 2017. *The Lancet*, 393(10184):1958–1972, 2019. doi: 10.1016/S0140-6736(19)30041-8.

[6] B.M. Bell, R. Alam, N. Alshurafa, E. Thomaz, A.S. Mondol, K. de la Haye, J.A. Stankovic, J. Lach, and D. Spruijt-Metz. Automatic, wearable-based, in-field eating detection approaches for public health research: a scoping review. *npj Digital Medicine*, 3(1):38, 2020. doi: 10.1038/s41746-020-0246-2.

[7] M.L. Neuhouser. The importance of healthy dietary patterns in chronic disease prevention. *Nutrition Research*, 70:3–6, 2019. doi: 10.1016/j.nutres.2018.06.002.

[8] F. Jannasch, J. Kröger, and M.B. Schulze. Dietary patterns and type 2 diabetes: a systematic literature review and meta-analysis of prospective studies. *The Journal of Nutrition*, 147(6):1174–1182, 2017. doi: 10.3945/jn.116.242552.

[9] N.D. Peterson, K.R. Middleton, L.M. Nackers, K.E. Medina, V.A. Milsom, and M.G. Perri. Dietary self-monitoring and long-term success with weight management. *Obesity*, 22(9):1962–1967, 2014. doi: 10.1002/oby.20807.

[10] N.A. Shaheen, A.A. Alqahtani, H. Assiri, R. Alkhodair, and M.A. Hussein. Public knowledge of dehydration and fluid intake practices: variation by participants' characteristics. *BMC Public Health*, 18(1):1346, 2018. doi: 10.1186/s12889-018-6252-5.

[11] J.A. Ship and D.J. Fischer. The relationship between dehydration and parotid salivary gland function in young and older healthy adults. *The Journals of Gerontology Series A: Biological Sciences and Medical Sciences*, 52(5):M310–M319, 1997. doi: 10.1093/gerona/52a.5.m310.

[12] L. Borghi, T. Meschi, T. Schianchi, A. Briganti, A. Guerra, F. Allegri, and A. Novarini. Urine volume: stone risk factor and preventive measure. *Nephron*, 81(Suppl. 1):31–37, 1998. doi: 10.1159/000046296.

[13] C.J. Edmonds, E.o Foglia, P. Booth, C.H.Y. Fu, and M. Gardner. Dehydration in older people: A systematic review of the effects of dehydration on health outcomes, healthcare costs and cognitive performance. *Archives of Gerontology and Geriatrics*, 95:104380, 2021. doi: 10.1016/j.archger.2021.104380.

[14] C.E. Ekpenyong and I.-A.M. Akpan. High prevalence and associated risk factors of dehydration among college students: implications for health and academic performance. *International Journal of Community Medicine And Public Health*, 4:1043–1055, 2017. doi: 10.18203/2394-6040.ijcmph20171322.

[15] E. Beard and R. West. Pilot study of the use of personal carbon monoxide monitoring to achieve radical smoking reduction. *Journal of Smoking Cessation*, 7(1):12–17, 2012. doi: 10.1017/jsc.2012.1.

[16] E.J. Perez-Stable, B.V. Marin, G. Marin, D.J. Brody, and N.L. Benowitz. Apparent underreporting of cigarette consumption among mexican american smokers. *American Journal of Public Health*, 80(9):1057–1061, 1990. doi: 10.2105/ajph.80.9.1057.

[17] J.K. Mandal, S. Misra, J.S. Banerjee, and S. Nayak (Eds.).Applications of Machine Intelligence in Engineering. *Proceedings of 2nd Global Conference on Artificial Intelligence and Applications (GCAIA, 2021)*, September 8-10, 2021, Jaipur, India, CRC Press, 2022. doi: 10.1201/9781003269793.

[18] S.-Y. Jeung and J.-W. Kwon. A robust multivariate time series classification approach based on topological data analysis for channel fault tolerance. *Sensors*, 25(9):2709, 2025. doi: 10.3390/s25092709.

[19] G.E. Carlsson. Topology and data. *Bulletin of the American Mathematical Society*, 46:255–308, 2009. doi: 10.1090/S0273-0979-09-01249-X.

[20] H. Edelsbrunner and J. Harer. Persistent homology-a survey. *Contemporary Mathematics*, 453(26):257–282, 2008. doi: 10.1090/conm/453/08802.

[21] Y. Skaf and R. Laubenbacher. Topological data analysis in biomedicine: A review. *Journal of Biomedical Informatics*, 130:104082, 2022. doi: 10.1016/j.jbi.2022.104082.

[22] J.A. Perea and J. Harer. Sliding windows and persistence: An application of topological methods to signal analysis. *Foundations of Computational Mathematics*, 15(3):799–838, 2015. doi: 10.1007/s10208-014-9206-z.

[23] J. Perea and C. Traile. Sliding windows and persistence. *The Journal of the Acoustical Society of America*, 141(5_Supplement):3585–3585, 2017. doi: 10.1121/1.4987655.

[24] P. Dłotko, W. Qiu, and S. Rudkin. Cyclicality, periodicity and the topology of time series. *arXiv preprint arXiv:1905.12118*, 2019. doi: 10.48550/arXiv.1905.12118.

[25] H. Poincaré. *Analysis situs*. Gauthier-Villars Paris, France, 1895.

[26] P. Bubenik. Statistical topological data analysis using persistence landscapes. *Journal of Machine Learning Research*, 16(1):77–102, 2015.

[27] R. Ghrist. Barcodes: The persistent topology of data. *Bulletin of the American Mathematical Society*, 45:61–75, 02 2008. doi: 10.1090/S0273-0979-07-01191-3.

[28] U. Bauer. Ripser: efficient computation of vietoris–rips persistence barcodes. *Journal of Applied and Computational Topology*, 5(3):391–423, 2021. doi: 10.1007/s41468-021-00071-5.

[29] D. Cohen-Steiner, H. Edelsbrunner, J. Harer, and Y. Mileyko. Lipschitz functions have L_p-stable persistence. *Foundations of Computational Mathematics*, 10(2):127–139, 2010. doi:10.1007/s10208-010-9060-6.

[30] M. Rucco, F. Castiglione, E. Merelli, and M. Pettini. Characterisation of the idiotypic immune network through persistent entropy. In *Proceedings of ECCS 2014: European conference on complex systems*, pages 117–128. Springer, 2016. doi: 10.1007/978-3-319-29228-1_11.

[31] O. Vasicek. A test for normality based on sample entropy. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 38(1):54–59, 1976. doi: 10.1002/9781119186229.ch35.

[32] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011. doi: 10.48550/arXiv.1201.0490.

[33] T. Akiba, S. Sano, T. Yanase, and M. Ohta, T.and Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2019. doi: 10.48550/arXiv.1907.10902.