

Explainable Machine Learning and Deep Learning Models for Predicting TAS2R-Bitter Molecule Interactions

Original

Explainable Machine Learning and Deep Learning Models for Predicting TAS2R-Bitter Molecule Interactions / Ferri, Francesco; Cannariato, Marco; Pallante, Lorenzo; Zizzi, Eric A.; Miceli, Marcello; Deriu, Marco A.. - In: JOURNAL OF MOLECULAR GRAPHICS & MODELLING. - ISSN 1093-3263. - 142:(2026). [10.1016/j.jmgm.2025.109187]

Availability:

This version is available at: 11583/3004052 since: 2025-10-15T10:45:45Z

Publisher:

Elsevier

Published


DOI:10.1016/j.jmgm.2025.109187

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)



Explainable machine learning and deep learning models for predicting TAS2R-bitter molecule interactions

Francesco Ferri^{a,b,1}, Marco Cannariato^{a,1}, Lorenzo Pallante^a, Eric A. Zizzi^a, Marcello Miceli^a, Marco A. Deriu^{a,*}

^a Politecnico di Torino, Polito^{BIO}MedLab, Department of Mechanical and Aerospace Engineering, Torino, 10129, Italy

^b TUM School of Life Sciences Weihenstephan, Technical University of Munich, Alte Akademie 8, 85354, Freising, Germany

ABSTRACT

This work aims to develop explainable models to predict the interactions between bitter molecules and TAS2Rs via traditional machine-learning and deep-learning methods starting from experimentally validated data. Bitterness is one of the five basic taste modalities that can be perceived by humans and other mammals. It is mediated by a family of G protein-coupled receptors (GPCRs), namely taste receptor type 2 (TAS2R) or bitter taste receptors. Furthermore, TAS2Rs participate in numerous functions beyond the gustatory system and have implications for various diseases due to their expression in various extra-oral tissues. For this reason, predicting the specific ligand-TAS2Rs interactions can be useful not only in the field of taste perception but also in the broader context of drug design. Considering that in-vitro screening of potential TAS2R ligands is expensive and time-consuming, machine learning (ML) and deep learning (DL) emerged as powerful tools to assist in the selection of ligands and targets for experimental studies and enhance our understanding of bitter receptor roles. In this context, ML and DL models developed in this work are both characterized by high performance and easy applicability. Furthermore, they can be synergistically integrated to enhance model explainability and facilitate the interpretation of results. Hence, the presented models promote a comprehensive understanding of the molecular characteristics of bitter compounds and the design of novel bitterants tailored to target specific TAS2Rs of interest.

1. Introduction

Taste perception is a crucial determinant of food intake and consumption patterns [1], with substantial consequences for human nutrition and health [2]. Several pathological conditions can impair the sense of taste and lead to alterations in quality of life and body weight regulation [3]. The molecular mechanisms of bitter taste perception involve a subfamily of 25 G protein-coupled receptors (GPCRs), called type 2 taste receptors (TAS2Rs) [4]. These proteins are also expressed in various extraoral tissues, where they may perform additional physiological functions, such as modulating inflammatory response, controlling upper respiratory immunity and others [5,6]. TAS2Rs exhibit a remarkable diversity and specificity in their ligand recognition and activation. Some TAS2Rs are defined as *promiscuous* as they can bind to multiple and structurally distinct bitter compounds, while others are highly *selective* and respond to only a few known ligands [7]. Conversely, some bitter compounds can activate several TAS2Rs, while others are specific for individual TAS2Rs. The chemical nature of bitter compounds is extremely heterogeneous and encompasses peptides, saponins, alkaloids, polyphenols, and salts [8]. In this context, understanding the

molecular interactions and predicting the specific association between bitter molecules and relative TAS2Rs could impact several fields of applications. For example, this line of research can: (i) help in the adherence to therapies based on bitter drugs [9]; (ii) improve the understanding regarding side effects or affect undiscovered biological pathways since TAS2Rs are present in extraoral tissues [10]; (iii) assist the design of alternative bitterants to improve the palatability of some healthy foods with high level of bitterness [11]. However, the current methods to identify the TAS2R targets for a given compound rely on laborious and costly in-vitro assays and, for this reason, the available information in the literature regarding protein-ligand interactions for bitter receptors is limited. Still, most of the data has been collected into a single database, called BitterDB [12].

To overcome the above-mentioned limitations, several computational approaches to predict the bitter taste or the interaction between bitter compounds and TAS2Rs have been proposed [13]. Computational methods offer several advantages, such as speed, low cost, and scalability, as well as the possibility of improvement as more experimental data become available. Advances in optimization algorithms have played a pivotal role in enhancing predictive modeling across a wide

* Corresponding author.

E-mail address: marco.deri@polito.it (M.A. Deriu).

¹ These authors contributed equally to this work.

range of scientific domains, including computational biology, medical diagnostics, and chemoinformatics [14–18]. These contributions have helped to refine both traditional and deep learning approaches, ultimately supporting more robust, generalizable, and interpretable models. Currently, many machine learning (ML) models have been proposed to accurately classify a compound as bitter or non-bitter, e.g. Ref. [19], BitterIntense [20], BitterCNN [21], and VirtuousSweetBitter [22]. Moreover BitterX [23], BitterSweet [24], and BitterMatch [25] can also predict the specific TAS2R target [26]. BitterX uses a Support Vector Machine (SVM) trained on a reduced and balanced dataset, whereas BitterMatch employs Gradient Boosting (GB) on Decision Trees (DTs) trained on data extracted from BitterDB [12]. On the other hand, although the BitterSweet webserver offers information on bitterant-TAS2R associations within its application, there is a lack of detail in both the original publication and the online materials regarding the specific development of this predictive model for association. While the above-mentioned algorithms rely on traditional ML, in recent years, Deep Learning (DL) models based on Neural Networks (NNs) have often offered increased performance. In particular, Graph Neural Networks (GNNs) are becoming widely used in computational chemistry since they are particularly suitable for the representation of molecular data [27]. In general, the common drawback of DL and many ML algorithms is their “black-box” nature which results in a complex explanation of the model’s predictions and feature importance. Therefore, several methods have been developed for both traditional ML algorithms and GNNs [28–30]. Among methods to explain traditional machine-learning methods, SHAP is a widely known model-agnostic method that can help in the understanding of how the features influence the prediction by elucidating toward which class the prediction is pushed [31]. On the other hand, tree-based algorithms can immediately output the feature importance but do not provide any further information. Specifically, Graph Neural Networks (GNNs) can directly represent molecular structures by encoding atoms and bonds as nodes and edges, respectively. This enables a more straightforward explanation of the model’s predictions on the molecular structures. For instance, both model-agnostic GNNExplainer [30] and Grad-CAM [32] methods unveil the influence of the single nodes and edges in the class decision, granting more visually impactful explanations.

In this work, we present two complementary models capable of predicting interactions between bitter compounds and TAS2Rs using traditional machine learning (ML) methods and graph neural networks (GNNs). These models could assist in identifying bitter compounds that selectively target specific bitter taste receptors and shed light on the molecular features underlying these interactions. Ultimately, this research could advance our understanding of the molecular features behind bitter taste perception and facilitate the development of tailored compounds designed to target specific TAS2Rs.

2. Material and methods

2.1. Dataset acquisition and preprocessing

The dataset comprises 338 bitter molecules with known positive and negative associations with 22 human bitter taste receptors (the remaining three receptors, namely TAS2R45, TAS2R48, and TAS2R60, are orphan and do not have any known agonists). Positive associations correspond to molecule-receptor couples that are known to interact, while negative interactions refer to molecules that do not bind to the associated receptor. Positive interactions are labelled as class “1” and negative as class “0”. Uniquely known and in-vitro verified interactions were considered, resulting in a total of 3964 paired associations. More in detail, 301 molecules (3204 known association with TAS2Rs) were taken from BitterMatch’s dataset [25], which is in turn derived from the BitterDB dataset [12]. The remaining 37 molecules (760 known association with TAS2Rs) were derived from recent literature [33–41]. The resultant dataset exhibits a remarkable imbalance, with the number of

binding instances approximately five times greater than the number of non-binding instances. In summary, the final dataset therefore considered 22 out of the 25 TAS2R receptors and consisted of 3964 pairs of bitterants-TAS2Rs known associations. The problem of predicting ligand-receptor association was reduced to a binary classification by defining the entries as bitterant-TAS2R pairs and the association as the only label. The resulting dataset is represented as a matrix where the number of rows equals the number of known associations. Each row contains information about the structure of the bitter compound, the associated receptor, and the target label, indicating the positive or negative association. Molecules are encoded as Canonical SMILES obtained either by BitterDB or PubChem [42]. Receptors are instead defined using the one-hot encoding strategy.

2.2. Interaction prediction using a traditional machine-learning (TML) approach

Molecules SMILES were standardized following the ChEMBL structure pipeline [43], as done in previous literature [22,44]. Fingerprints and physicochemical descriptors were used as features for the bitter molecules. Morgan fingerprints (number of bits = 1024, radius = 2) were computed using RDKit python package starting from the previously sanitized SMILES. Molecular descriptors were calculated with the Mordred Python library [45]. All descriptors having more than 90 % correlation with other descriptors were removed. All non-binary data was normalized with Min-Max normalization [46] (see also Supplementary Information for further details on the normalization procedure).

Several traditional machine learning algorithms, namely Gaussian Naive Bayes (GaussianNB), Logistic Regression (LR), K-Neighbors, Support Vector Machines (SVM), Random Forest (RF), and Gradient Boosting on Decision Trees (GB on DTs), were first compared. The model demonstrating the highest performance (GB on DTs) was subsequently chosen as the optimal model for further analyses. This algorithm was also used in similar models [25] and demonstrated remarkable resilience towards overfitting and biases, which are essential when dealing with imbalanced data [47]. CatBoost [48], a high-performance open-source library for gradient boosting on decision trees, was employed. CatBoost uses a combination of ordered boosting, random permutations, and gradient-based optimization to achieve high performance on any dataset. The comprehensive workflow for the TML approach is shown in Fig. 1.

The molecular dataset first underwent a chemical space analysis. Employing a clustering algorithm on this matrix, we grouped molecules based on their structural features. We used agglomerative clustering (complete linkage algorithm) to divide data into n clusters with similar variances. To define the optimal n number of clusters the Silhouette score analysis over different n values was carried out. The parameter chosen as the distance (d) to divide the chemical space into clusters with Agglomerative clustering was the Tanimoto [49], computed from the precomputed Morgan fingerprints. The molecules were randomly selected from each cluster with an 80:20 ratio, to obtain a split stratified on the cluster label, ensuring each cluster was well represented in both sets. The total number of ligand-based features was 2824. Dimensionality reduction of the dataset is, therefore, necessary to keep only the most informative features and simplify the model explanation. Feature selection was performed using two methods, i.e. (i) the “noisy” feature selection approach and (ii) Sequential Feature Selection (SFS). The noisy feature selection method is based on previous literature [50] and adopts an iterative selection technique by incorporating a random column into the dataset. Termed “noisy”, this column is populated with pseudo-random numbers ranging from 0 to 1. Following the training phase of the tree-based classifier, the Gini importance for each feature was computed. Features with lower importance than the noisy feature were then systematically excluded from the dataset until only features more informative than the noisy column remained. The SFS is a greedy

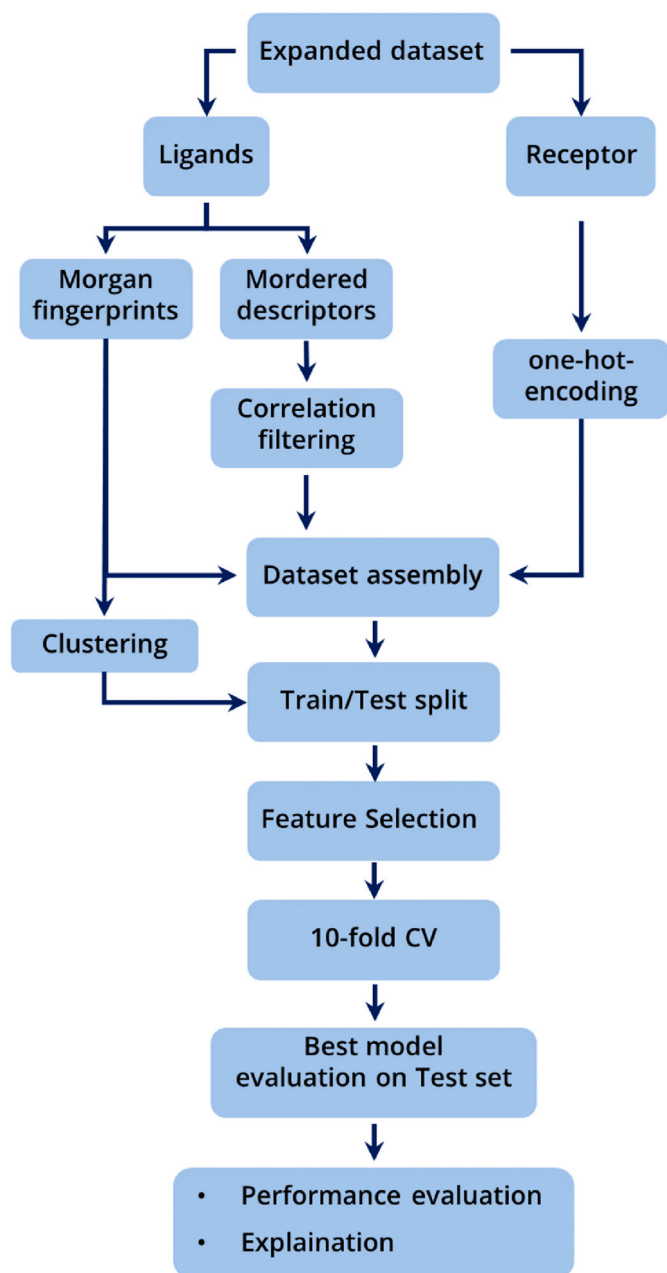


Fig. 1. Traditional Machine-Learning (TML) workflow.

algorithm, that selects a subset of features in a forward or backward direction. The selection or removal of a feature depends on the cross-validation (CV) score of an estimator trained on the current feature subset. Forward-SFS begins with no features and gradually adds the best feature that maximizes the CV score until it reaches the desired number of features. Similarly, the Backwards-SFS starts with all features and progressively removes the worst feature that minimizes the CV score [51]. Scikit-learn was employed to perform backward-SFS considering the average precision as the 5-fold CV score, starting from the 150 most important features according to the CatBoostClassifier tree-based importance metric. The final number of features was decided a posteriori.

Finally, the classification task was then carried out by training the CatBoostClassifier on the training set, using a 10-fold CV, and evaluating the model's performance on the test set.

2.3. Interaction prediction using graph convolutional neural networks

Graph Convolutional Networks (GCNs) extend the neural network models to work on arbitrarily structured graphs using graph convolutions and spectral filters [52]. A GCN is an approach for semi-supervised learning on graph-structured data. It is based on an efficient variant of convolutional neural networks which operate directly on graphs. In particular, spatial-based methods define graph convolutions by using the node's spatial relations similarly to the convolutional operator of a CNN on an image. The proposed model is built on PyTorch and PyTorch Geometric [53,54].

Referring to the Graph as $G = (V, E)$, V is the set of nodes and E is the set of edges. Suppose v_i is a node in V and $e_{ij} = (v_i, v_j)$ is an edge that originates from v_i and terminates at v_j . The graph is therefore defined by n nodes, m edges, and d, b , and c as dimensions of the node, hidden node, and edge feature vectors, respectively [55]. The set of nodes that are adjacent to a node v is denoted by $N(v) = \{u \in V | (v, u) \in E\}$. The graph has an adjacency matrix A of size $n \times m$, such that $A_{ij} = 1$ if and only if e_{ij} belongs to E , and $A_{ij} = 0$ otherwise. The graph also has node features X , which is a $n \times d$ matrix of node feature vectors, each of which has d dimensions and corresponds to a node v , i.e., x_v . Additionally, the graph may have an edge feature matrix X^e , of dimension $m \times c$, composed by edge feature vectors of c dimensions and represents an edge (v, u) , i.e. $x_{v,u}^e$. A molecule is defined as an undirected graph in which the adjacency matrix is always symmetric.

Molecules, stored as standardized SMILES, were converted into molecular graphs using NetworkX [56]. The node and edge features were selected according to previous literature [57–59] and are listed in Table S1. As in the TML approach, the dataset was clustered based on the Tanimoto similarity before splitting each cluster into training and test tests (80:20) by stratifying over the class labels. The developed model takes as input a batch of graphs, each with node features, edge features, and the one-hot encoded receptor, and predicts the positive or negative association with the specified receptor. The overall workflow for the GCN approach is shown in Fig. 2.

The GCN model consists of four main components: the graph convolutional layers, the batch normalization layers, the dropout layers, and the fully connected layers. The graph convolutional layers use the GATv2Conv module, which implements the Graph Attention Network (GAT) variant proposed by previous literature [60]. The GAT module computes the node embeddings by applying a self-attention mechanism over the node features and the edge features. The model uses two graph convolutional layers, with 32 and 8 output channels, respectively. Batch normalization layers are placed after both convolutional layers. Graph embeddings are then mapped using four fully connected layers, with 32, 16, 8, and 4 output units, respectively. The model uses two dropout layers, one with a probability of 0.1 and one with a probability of 0.2. The first dropout layer is applied to the input of the fully connected layers, and the second dropout layer is applied to the output of the last fully connected layer. The first fully connected layer takes as input the concatenation of the graph embeddings, obtained by applying global mean pooling to the node embeddings, and the receptor features. The last fully connected layer is followed by an output layer, which uses a linear transformation to produce two outputs. The model makes use of the ReLU activation function for the hidden units and of a sigmoid activation function for the node embeddings. The model returns the output of the output layer, which can be interpreted as the probability of each class. The architecture of the GCN model is schematically represented in Fig. 3.

2.4. Explainability

Regarding the TML, CatBoost evaluates the relevance of each input feature by using the individual importance values. These values reflect the average variation in the prediction caused by the modification of the

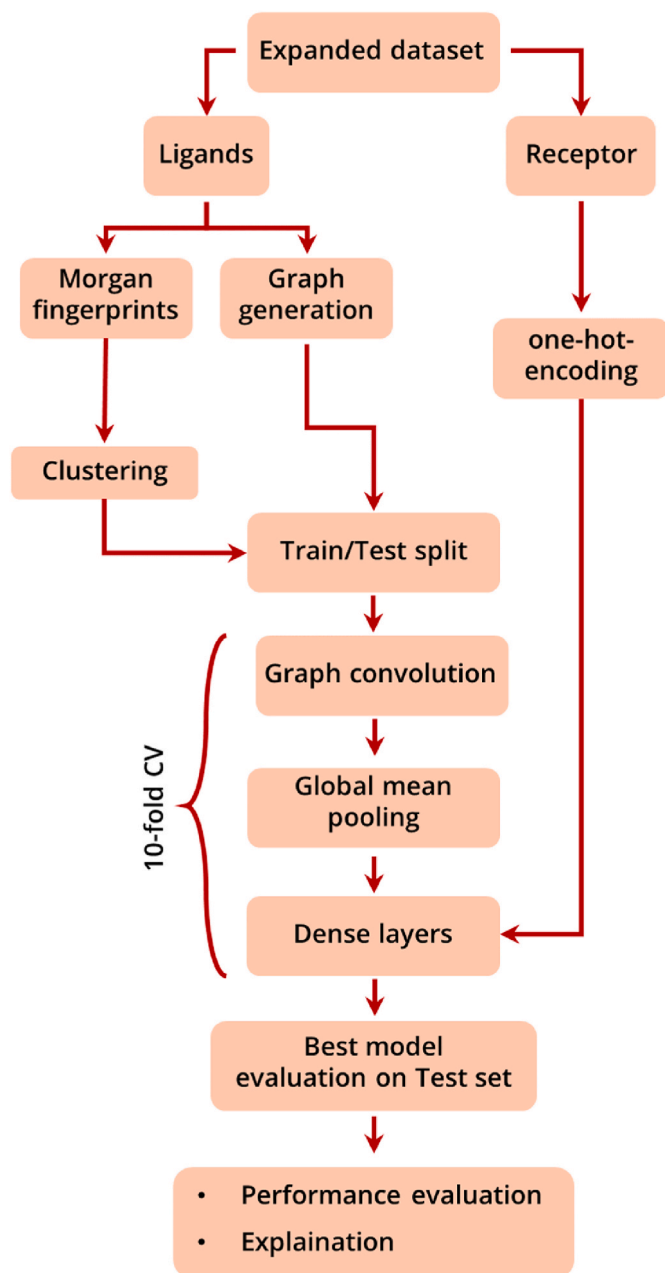


Fig. 2. GCN framework workflow.

feature value. Moreover, we employed SHAP (SHapley Additive exPlanations) to underscore the specific feature importance [31]. SHAP values are a way to interpret the output of any machine learning model, based on the principles of game theory. SHAP values exhibit consistency, indicating that features deemed unequivocally more important will consistently yield higher SHAP values. The employed SHAP library uses a Tree-based model as the local explanation method for trees [28], allowing the calculation of optimal local explanations.

The explainability of the GCN model was achieved through GNNExplainer [30] and Grad-CAM [32] methods. GNNExplainer is a model-agnostic approach for providing interpretable explanations for predictions of any GNN-based model on any graph-based machine learning task [30], allowing for single-instance explanations. At the same time, the unfaithfulness of the explanation was evaluated using the graph explanation faithfulness (GEF) score [61], calculated as:

$$GEF(y, \hat{y}) = 1 - e^{-KL(y|\hat{y})}$$

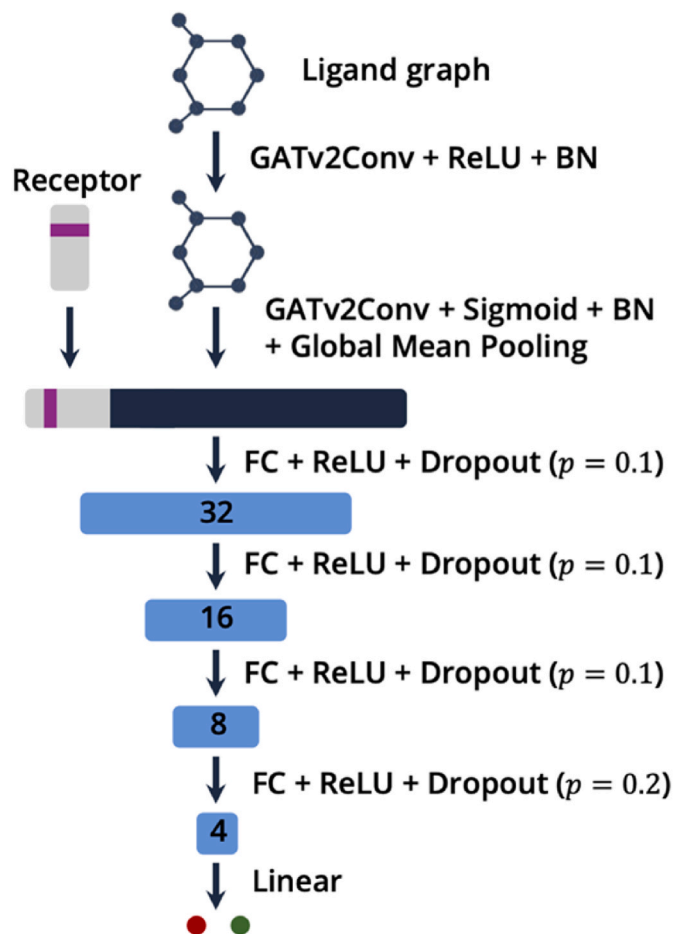


Fig. 3. Schematic representation of the GCN model architecture.

where y and \hat{y} refer to the output probability vector obtained from the original graph and masked subgraph, respectively, and KL is the Kullback-Leibler divergence score. Therefore, the GEF score ranges between 0 and 1, with values near 0 indicating excellent prediction faithfulness and values close to 1 indicating very poor faithfulness. Typically, faithfulness scores higher than 0.5 are considered indicative of untrustworthy explanations.

On the other hand, Grad-CAM, originally developed to identify salient regions in image classification problems [32], considers the gradient of the output with respect to the last convolutional layer. In this work, a generalization to graphs in the form of Unsigned Grad-CAM (UGrad-CAM) was employed to identify positive and negative contributions from each node [29]. The implementation of both Grad-CAM and UGrad-CAM was adapted from publicly available code repositories (<https://github.com/ndey96/GCNN-Explainability>).

3. Results

3.1. Interaction prediction using a traditional machine-learning (TML) approach

Initially, we compared traditional machine learning algorithms, including Gaussian Naive Bayes (GaussianNB), Logistic Regression (LR), K-Neighbors, Support Vector Machines (SVM), Random Forest (RF), and Gradient Boosting on Decision Trees (GB on DTs). GB on DT achieved better performances in terms of ROC (Receiver Operating Characteristics) AUC (Area Under the Curve) (see also Fig. S1) and was therefore selected as the best model for the traditional machine learning approach. The CatBoostClassifier hyperparameters were tuned on the

training set and are detailed in Table S2. Then, two feature selection methods, namely the “noisy” and the Backward-sequential feature selection (SFS) methods, were compared to select only the most informative ligand-based features, as also detailed in the Materials and Methods section. 28 ligand features were selected using the “noisy” method, whereas 17 through the SFS approach (Fig. S2). The two feature selection methods achieved similar performance in terms of ROC and Precision-Recall (PR) AUC on the test set (Fig. S3). It is interesting to note that both methods ultimately selected only Mordred descriptors, while no ligand fingerprints were retained in the final feature set. Given the higher reproducibility and the lower number of selected features, the SFS method was selected as the most suitable for the feature selection task. Therefore, from now on, the GB on DT using the SFS method be referred to as the *TML model*. The relative ROC and PR curves on the test set, reported in Fig. 4, were characterized by AUC values of 0.92 and 0.75, respectively.

The tree-based importance of the employed features was evaluated (Fig. 5). The importance of the Mordred descriptors is in general the highest, with the only exception of the association with TAS2R14 and TAS2R46, which are the first and third most important features according to the model. Notably, these TAS2Rs are also the two most promiscuous receptors.

The feature importance of the TML model was also assessed through the SHAP method (Fig. 6). Comparing the SHAP values for the one-hot encoded receptor features (Fig. 6B) to the number of agonists of each TAS2R (Fig. 6A), it could be observed that associations with the most promiscuous receptors tend to bias the prediction toward class 1, whereas connections with more selective receptors lean towards the class 0. Moreover, SHAP was employed to evaluate how specific features influence individual predictions. For example, the explanation for the strychnine-TAS2R46 pair (positive association) and the strychnine-TAS2R1 pair (negative association) are shown in Fig. 6C and D. It is noteworthy that the association feature with the respective receptor typically holds a notable influence on the prediction outcome. However, in certain instances, such as observed with the strychnine-TAS2R1 pair, other features may also remarkably impact the final predicted association.

3.2. Interaction prediction using graph neural network approach

Similarly to the TML model, the ROC and PR curves for the GCN classifier were generated by conducting a 10-fold CV and selecting the model with the highest performance in the best fold to evaluate it on the Test Set (Fig. 7). The GCN model exhibited a ROC AUC of 0.88, demonstrating performance on the Test set comparable to that observed during the CV. Regarding the PR curve, the model was characterized by an AUC of 0.67.

We then employed the GNNExplainer [30] and Grad-CAM [32] methods to develop single-instance explanations for the GCN model predictions. Considering the case of the strychnine-TAS2R46 association (class 1) as an example, the GNNExplainer made it possible to inspect the importance of node features for the prediction (Fig. 8A) as well as of each bond within the molecule (Fig. 8B). Notably, the three most important node features include the atom’s partial charge and partition coefficient, which can be related to the hydrophilicity of the molecule. Moreover, the bonds formed by a tertiary amine, which was experimentally found to be involved in an interaction with TAS2R46, are among the most relevant for the prediction. On the other hand, UGrad-CAM analysis was used to inspect the contribution of each node in the prediction toward the two classes. Such information, for the strychnine-TAS2R46 pair, is shown as a heatmap in Fig. 8C. Interestingly, the region involving the aforementioned tertiary amine holds a great contribution toward class 1. On the other hand, the aromatic ring of strychnine represents the main contribution towards class 0. This explainability approach can be also useful to inspect the possible effect of structural alterations of the molecule on the output prediction. For




Fig. 4. (A) TML’s ROC curves for the validation (green, mean and standard deviation during the 10-fold CV), and test (red) sets. (B) PR curve on the test set. (For interpretation of the references to colour in this figure legend, the reader is referred to the Web version of this article.)

instance, removing two carbon atoms from strychnine near the above-mentioned tertiary amine changes the pattern of importance values (Fig. 8D). Such modification results in a contribution toward class 0 for this amine, which becomes secondary, together with a reduced influence for other regions of the molecule. As a result, the output probability of the model was reduced.

Finally, the performance of the TML and GCN models on the test set was compared (Table 1). Despite GCN achieving higher recall and F2 scores on class 1 compared to the TML model, the TML model is characterized, on average, by higher values of the considered metrics. In




Fig. 5. Tree-based feature importance of the TML model. The 22 receptor association features are displayed in orange, while the 17 ligand descriptors selected by the SFS method are shown in blue. (For interpretation of the references to colour in this figure legend, the reader is referred to the Web version of this article.)

particular, the TML precision is remarkably higher than the GCN precision for class 1. Therefore, regarding the under-represented class, TML seemed to prioritize the precision while GCN the recall.

4. Discussion

The present work aims to develop a machine-learning-based model to predict which TAS2R receptor is targeted by a specific bitter molecule based on its molecular structure. To achieve this goal, two complementary approaches using traditional machine learning methods and graph-based neural networks have been developed.

The dataset consists of 343 bitter molecules for which the positive (binders) or negative (non-binders) association with the human TAS2R bitter taste receptors is experimentally known. The resulting dataset is therefore composed of 3970 bitterant-TAS2R pairs with a relative label corresponding to a positive (class 1) or negative (class 0) interaction. A major obstacle to the present work is the paucity, diversity and unbalancing of the available data on TAS2R-ligand interactions. The dataset contains only bitter molecules, limiting the applicability of the models only to this specific taste. Regarding the applicability domain (AD) of the models, we evaluated the reliability of the model predictions based on

the similarity between the tested compounds and the chemicals used during the training phase. We used an average-similarity approach already employed in previous literature in the taste prediction field [19, 62] and in our previous works [22,44] (see also Supplementary Information and Fig. S4). The AD check is performed every time before running the model to assess the reliability of the prediction and the output of the AD control is given to the user. Future experimental studies elucidating the interactions between other bitter compounds or other chemicals with TAS2Rs would considerably enhance the performance of the models and broaden their chemical applicability domain. Moreover, the models could be improved by adding features related to the three-dimensional structure of the bitter taste receptors that are known to drive the ligand binding and recognition, such as the volume of the binding pocket, the Solvent Accessible Surface Area (Area), the radius of gyration, and many others highlighted in previous literature [25]. Indeed, the accurate experimental or *in silico* determination of the molecular structures of bitter receptors, akin to GPCRs in general, remains a complex and ongoing challenge in this area of research.

After comparing several traditional machine learning algorithms, Gradient Boosting on Decision Trees (GB on DTs) was chosen as the best architecture for the present problem (see also Fig. S1). Interestingly, the RandomForest achieved similar performance to the GB on DTs, confirming the soundness of tree-based algorithms for solving this specific prediction problem. It is noteworthy that even the most recent comparable work in the literature, such as BitterMatch, relies on a tree-based method (XGBoost) [25]. To reduce the number of ligand-based features, the GB on DT was then tested with two feature selection methods: (i) the “noisy” feature selection method, which preserves only features with importance greater than a random-generated feature; and (ii) the Backward Sequential Feature Selection (SFS) method, which evaluate the performance of the model by iteratively removing a feature from the dataset. The performance with the two feature selection methods showed similar results both on ROC and PR curves (see also Fig. S3). For this reason, despite its higher computational cost, the Backward-SFS was preferred to the “noisy” feature selection method since it is completely reproducible and since it allowed to select only 17 features, compared to the 28 selected by the “noisy” method.

A tree-based algorithm also facilitated the assessment of feature importance, as shown in Fig. 5. The most important feature is the association with receptor 14, noticeably the most promiscuous receptor with the highest number of known agonists, followed by receptor 46, which is the second most numerous in terms of known binders. As clearly shown in the bar plot, the more selective is the receptor, the less the importance value of their association. On the other hand, regarding the ligand-based descriptors, the SFS method selected 17 features, comprising only Mordred molecular descriptors and no ligand fingerprints. More in detail, the most informative features are: GATS1i (geary coefficient of lag 1 weighted by ionization potential), ATSC4d (centered moreau-broto autocorrelation of lag 4 weighted by sigma electrons), Xpc-5dv (5-ordered Chi path-cluster weighted by valence electrons), VR2_Dzm (normalized Randic-like eigenvector-based index from Barysz matrix weighted by mass), SIC4 (5-ordered structural information content), AXp-1dv (1-ordered averaged Chi path weighted by valence electrons), ATSC7are (centered moreau-broto autocorrelation of lag 7 weighted by allred-rocw EN), MATS3s (morán coefficient of lag 3 weighted by intrinsic state), AATS1i (averaged moreau-broto autocorrelation of lag 1 weighted by ionization potential), Mv (mean of constitutional weighted by vdW volume), GATS3v (geary coefficient of lag 3 weighted by vdW volume), AMID_C (averaged molecular ID on C atoms), ATSC4s (centered moreau-broto autocorrelation of lag 4 weighted by intrinsic state), ATSC4pe (centered moreau-broto autocorrelation of lag 4 weighted by pauling EN), ATSC3i (centered moreau-broto autocorrelation of lag 3 weighted by ionization potential), EState_VSA7 (EState VSA Descriptor 7 (1.81 ≤ x < 2.05)), and GGI7 (7-ordered raw topological charge). The comprehensive compilation of employed Mordred descriptors can be accessed at <https://mordred-des>

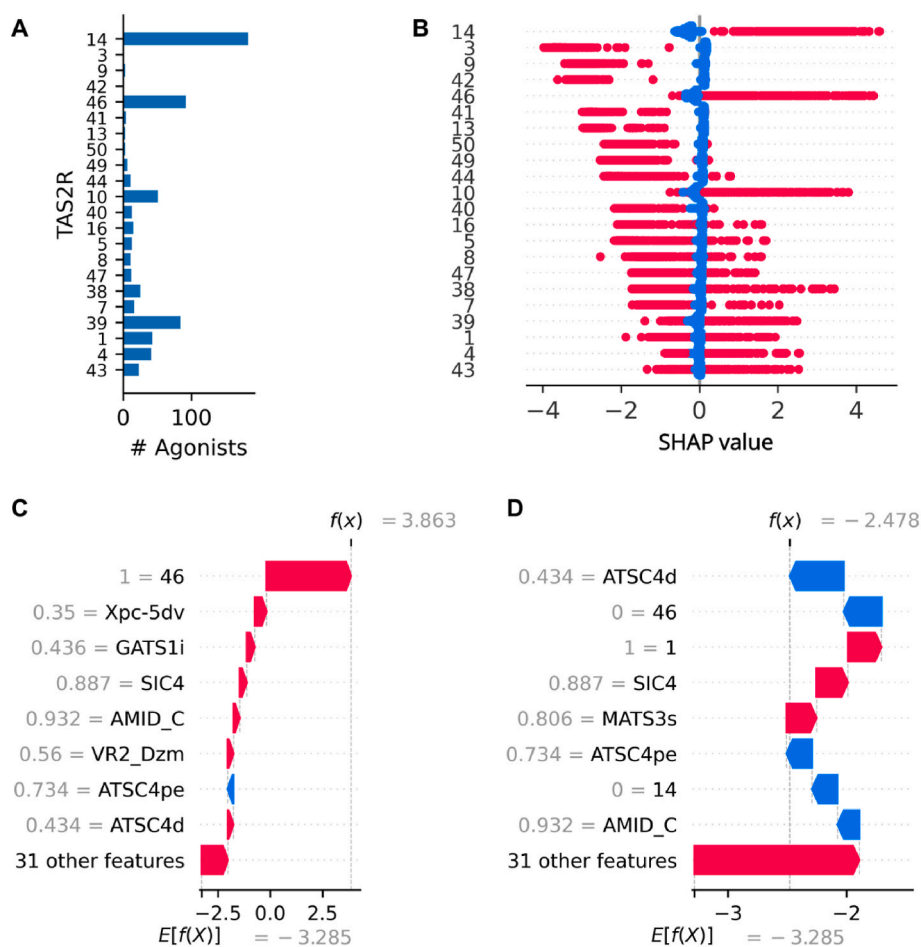


Fig. 6. (A) Number of agonists per TAS2R and (B) relative SHAP values for receptor associations. (C, D) SHAP waterfall for strychnine-TAS2R46 (positive association) (C) and strychnine-TAS2R1 (negative association) (D) pairs. $E[f(X)]$ is the average raw prediction (log odds) across the dataset and $f(x)$ is the prediction for the specific pair (log odds). Blue bars correspond to negative contributions to the prediction, while red bars to positive contributions. Features and their value for the specific pair are reported on the vertical axis. (For interpretation of the references to colour in this figure legend, the reader is referred to the Web version of this article.)

riptor.github.io/documentation/master/descriptors.html. The descriptors with higher occurrence fall into the categories of autocorrelation and topological descriptors. Autocorrelation descriptors (GATS1i, ATSC4d, ATSC7are, ATSC4s, ATSC4pe, ATSC3i) capture spatial relationships between atoms or properties within a molecule, while topological (Xpc-5dv, MATS3s, GGI7) descriptors focus on the molecular graph's structure. Despite the substantial reduction in the number of features, comprehending the chemical and physical properties of tastants based on the 17 most important features remains challenging. To improve the model's explainability, future studies should prioritize the utilization of simpler descriptors or the development of specific methodologies to intuitively relate the molecular descriptors with the relative structural features or functional groups. In this context, the explainability obtained for the GCN model is considerably easier to understand and to directly correlate molecular features to the overall prediction.

Moreover, SHAP values highlighted not only the importance of the single features on the final prediction but also their specific influence on the final prediction. Firstly, Fig. 6A and B shows that pairs involving more promiscuous receptors (like TAS2R14, 46, and 39) pull the prediction toward the positive class. On the other hand, being associated with a selective receptor favours the negative class. SHAP feature importance also allows to obtain a local interpretation of the model, analysing a specific bitterant-TAS2R association. For example, Fig. 6C shows how the feature's value influences the prediction of the strychnine-TAS2R46 pair (known association: positive – class 1,

prediction: class 1), while Fig. 6D for strychnine-TAS2R1 (known association: unknown, prediction: class 0). In the case of the strychnine-TAS2R46 pair, the feature associated with TAS2R46 exerts a substantial influence, driving the prediction towards the positive class. This finding aligns with expectations, as TAS2R46 is recognised as one of the most promiscuous bitter taste receptors, and the dataset consequently incorporates numerous instances of interacting compounds. Conversely, for the strychnine-TAS2R1 interaction, the most influential feature is a ligand-based descriptor (ATSC4d), suggesting that the molecular structure of the ligand is the primary factor driving the prediction in this instance. It is important to underline that due to the dataset's imbalance, with approximately five times as many negative samples as positive ones, the overall prediction tends to skew towards the negative class. Specifically, lower negative SHAP values exert a stronger influence on pushing predictions towards the negative class, while lower positive SHAP values primarily serve to counterbalance the inherent negativity of the initial prediction. Given that, the examples in Fig. 6 show that in most cases the sum of the features pushes toward a class, while few Mordred descriptors and receptor associations could remarkably alter the result of the prediction.

Regarding the implementation of the Graph Convolutional Neural Network (GCN), we converted the dataset of molecules into a graph-based representation using NetworkX. Each node and edge of the graphs are described by a specific set of features as listed in Table S1. The developed GCN model achieved on the test set ROC AUC and PR




Fig. 7. (A) GCN's ROC curves for the validation (green, mean and standard deviation during the cross-validation), and test (red) sets. (B) PR curve on the test set. (For interpretation of the references to colour in this figure legend, the reader is referred to the Web version of this article.)

AUC scores of 0.88 and 0.67, respectively (Fig. 7). In terms of AUC, the TML outperformed the GCN model reaching 0.92 and 0.75 for ROC AUC and PR AUC scores, respectively (Table 1). It is worth noticing that TML and GCN have comparable performance in the negative samples (class 0), whereas the precision of the TML in the positive samples (class 1) is remarkably higher (TML: 0.78; GCN: 0.62). This discrepancy in performance could be attributed to the dataset's imbalance, where there are nearly five times more negative pairs compared to positive samples. This imbalance likely affects the performance, resulting in lower precision for the positive class compared to the negative class. This issue is

exacerbated for the GCN model due to its increased sensitivity to limited samples in the training set.

The GNNExplainer [30] and Grad-CAM [32] methods were employed to explain the GCN model. GNNExplainer offers insights at a graph level, considering the entirety of the molecule, while Grad-CAM can highlight the node importance and provide information on its impact towards positive or negative class. The GNNExplainer on a graph level can easily be compared with the Tree-based feature importances. Noticeably, features of the molecule directly linked to specific types of interactions (electrostatics, polar interactions, hydrogen bonds, etc.) with the receptor are identified as the most important by the GNNExplainer. In particular, for the strychnine molecule (Fig. 8A), the most important node features include the atom partial charge (Gasteiger Charge) and partition coefficient (logP), which can be related to the hydrophilicity of the molecule, a crucial factor for the ligand-target interaction. Moreover, edge importance emphasises the structural motifs that affect the model's output (Fig. 8B). On the contrary, Grad-CAM and UGrad-CAM can generate heatmaps of gradient-weighted class activations, offering a robust visual explanation of the atoms influencing the prediction and indicating which class each node is favouring. Fig. 8C shows UGrad-CAM applied on the Strychnine-TAS2R46 pair. This association was chosen specifically because of recent experimental evidence regarding strychnine's binding to TAS2R46 [63]. The authors suggest that the binding occurs thanks to a π - π interaction between residue W88 and strychnine's benzene ring, and with the formation of a hydrogen bond between residue E265 and one of the two tertiary amines of strychnine. By looking at the UGrad-CAM plot (Fig. 8C), it is interesting to notice that the same tertiary amine involved in the binding is one of the most activated atoms towards class 1 prediction. Interestingly, the GNNExplainer analysis also consistently identified edges connected to this tertiary amine as crucial for the final prediction (Fig. 8B). In summary, the explainability of the GCN model appears particularly promising in identifying the structural molecular features underlying the ligand-receptor association and facilitating the explanation of the model.

The proposed models have been also compared with previous examples in the literature. At present, as far as the authors know, only three methods, i.e. BitterX [23], BitterSweet [24], and BitterMatch [25], have been developed to predict the interaction between bitterants and TAS2Rs. We exclusively compared our model with BitterMatch as it is the most recently developed tool and has demonstrated the best performance. Additionally, while the BitterSweet web server provides predictions on bitterant-TAS2R associations, no information is available regarding the specific development of the BitterSweet model for these associations in either the original publication or the online materials. Moreover, BitterX was not considered for a fair comparison because it was trained on a profoundly different dataset (540 bitterants with 260 positive and 260 negative bitterant-TAS2R interactions). To effectively compare BitterMatch with the present work, the BitterMatch model from the official GitHub repository (<https://github.com/YuliSI/BitterMatch>) was trained only on human data (*BM Human-Only*) by removing data relative to murine receptors. The comparison of the AUC and PR curves for the TML, GCN and BM-Human-Only models are reported in Fig. S5, whereas all performance metrics are summarised in Table S3. As indicated in the comparison between the TML and GCN models, performance metrics for class 0 are higher and almost identical, while performance for class 1 is notably lower and with some differences for the three compared models. Consistent with our observations in the comparison between the TML and GCN models, the performance on class 0 (negative associations) is again notably higher than on class 1 (positive associations) for the BitterMatch (BM) model. This disparity can again be attributed to the dataset being unbalanced towards the negative class. Overall, the three models achieved similar PR AUC scores on their test sets, and the other metrics indicated that the three models are nearly equivalent, with BitterMatch achieving slightly less performance in Recall, F1, and F2 in class 1. However, it is worth mentioning

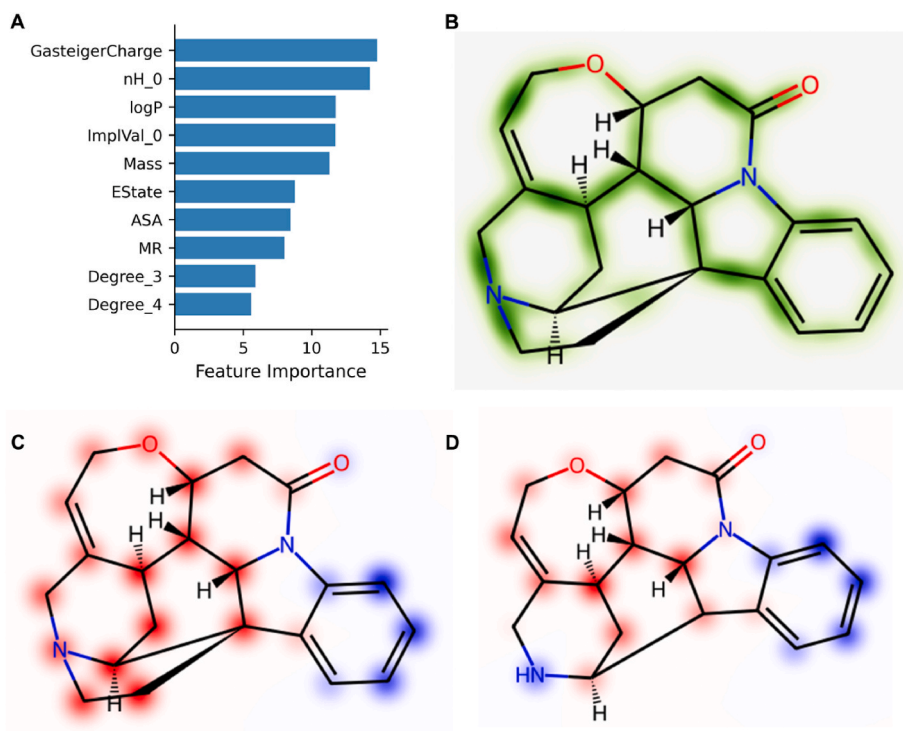


Fig. 8. GCN model explainability for the strychnine-TAS2R46 pair (positive association) using GNNExplainer (A, B) and UGrad-CAM (C, D). (A) Node feature importances for the 10 most important features for the prediction and (B) heatmap of edge importance, where the darkness of the green is proportional to the importance of the edge. (C) UGrad-CAM heatmap where red nodes correspond to contribution towards class 1 and blue towards class 0. (D) UGrad-CAM heatmap of a modified strychnine molecule paired to the TAS2R46 receptor, where red nodes correspond to contribution towards the predicted class (1), and blue towards the opposite class (0). (For interpretation of the references to colour in this figure legend, the reader is referred to the Web version of this article.)

Table 1

Comparison of TML and GCN performance metrics on the Test set. Values in bold represent the highest of the metrics compared.

	ROC AUC	PR AUC	Class	Precision	Recall	F1	F2
TML	0.92	0.75	0	0.93	0.97	0.95	0.96
			1	0.78	0.60	0.68	0.63
GCN	0.88	0.67	0	0.94	0.92	0.93	0.93
			1	0.62	0.67	0.64	0.66

that, unlike BitterMatch, the present models are capable of predicting the class of any query molecule within the model's applicability domain using only its SMILES representation. This expands their usage to a wider audience and paves the way toward the possibility of using these methodologies in emerging fields of interest, such as precision nutrition or nutraceutical development. Additional evidence of the models' robustness is demonstrated by the results obtained from an external dataset made available during the article's writing (see also Fig. S6 and Table S4), showcasing the effectiveness of the methods.

In summary, when considering only performances on prediction, the TML model achieves a better Precision score when compared to the GCN model, also offering a better alternative to other state-of-the-art models like BitterMatch in terms of Recall, F1 and F2 scores. On the other hand, GCN possesses the ability to generate more impactful visual explanations, which could be easily interpreted and exploited for subsequent tasks. Hence, we regard the Traditional Machine Learning (TML) and Graph Convolutional Network (GCN) models presented in this study as valuable complementary tools, offering robust predictions of TAS2Rs-bitterants interactions along with insightful visual explanations that connect molecular structures with their targeting capabilities.

5. Conclusion

This work presents a novel approach to predicting the interactions between bitter taste receptors and their ligands, using both a Traditional Machine-Learning (TML) approach and Graph Convolutional Neural Networks (GCN). Both models were also designed to be explainable, either by using interpretable algorithms, such as decision trees, or by applying custom methods, such as SHAP, GNNExplainer, and Grad-CAM, to reveal the most important features and molecular motifs underlying the predictions. The results indicated that both the TML and GCN models achieved similar and satisfactory predictive performance. Compared to previous works, the developed models offer ease of use, applicability to new molecules within the model's applicability domain, and interpretability, providing valuable insights into the associations between molecular features and their corresponding bitter taste receptor targets. While the TML model demonstrated slightly superior reliability in terms of performance metrics, the GCN model offers explanations through visually rich representations directly of the molecular structure of the compounds under investigation, thereby enhancing the interpretability of model predictions.

In a broader context, the *in silico* identification of promising compounds capable of targeting TAS2Rs holds significant potential across multiple fields. These methodologies can aid in the development of novel bitter modulators and enhancers, exerting a substantial impact not only in the food industry but also in the pharmaceutical sector, where the bitter taste of medications often poses challenges to effective administration for children or the elderly. Furthermore, since certain bitter receptors are also found in extra-oral tissues and are implicated in various diseases such as obesity, diabetes, asthma, and cancer, the development of these machine-learning models represents a promising strategy for identifying specific ligands capable of targeting these proteins. Such models have the potential to overcome some of the limitations associated with traditional experimental approaches, including

high cost, low throughput, and ethical issues.

Funding

The work was developed within the VIRTUOUS (Grant Agreement No 872181) and GALATEA (Grant Agreement No. 101183057) projects, funded by the European Union under the Marie Skłodowska-Curie Actions.

CRediT authorship contribution statement

Francesco Ferri: Validation, Data curation, Writing – review & editing, Software, Conceptualization, Visualization, Formal analysis, Writing – original draft, Methodology. **Marco Cannariato:** Writing – review & editing, Software, Conceptualization, Writing – original draft, Methodology, Validation, Data curation, Visualization, Formal analysis. **Lorenzo Pallante:** Writing – original draft, Visualization, Writing – review & editing, Conceptualization, Data curation. **Eric A. Zizzi:** Writing – review & editing, Supervision, Conceptualization. **Marcello Miceli:** Writing – review & editing, Supervision. **Marco A. Deriu:** Resources, Writing – review & editing, Project administration, Supervision, Conceptualization, Validation, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The authors acknowledge the CINECA award under the ISCR initiative for the availability of high-performance computing resources and support.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.jmngm.2025.109187>.

Data availability

All codes and data needed to replicate the results of the present work are accessible at <https://github.com/francescofers/TAS2RsPredictor>.

References

- J.I. Glendinning, Is the bitter rejection response always adaptive? *Physiol. Behav.* 56 (6) (1994) 1217–1227, [https://doi.org/10.1016/0031-9384\(94\)90369-7](https://doi.org/10.1016/0031-9384(94)90369-7).
- A. Shahbandi, E. Choo, R. Dando, Receptor Regulation in Taste: Can Diet Influence How We Perceive Foods? *J.* 1, 2018, pp. 106–115, https://doi.org/10.3390/j1010011_1.
- D. Risso, D. Drayna, G. Morini, Alteration, reduction and taste loss: main causes and potential implications on dietary habits, *Nutrients* 12 (11) (2020) 11, <https://doi.org/10.3390/nu12113284>.
- L. Pallante, M. Malavolta, G. Grasso, A. Korfiati, S. Mavroudi, B. Mavkov, A. Kalogeras, C. Alexakos, V. Martos, D. Amoroso, G. di Benedetto, D. Piga, K. Theofilatos, M.A. Deriu, On the human taste perception: molecular-Level understanding empowered by computational methods, *Trends Food Sci. Technol.* 116 (2021) 445–459, <https://doi.org/10.1016/j.tifs.2021.07.013>.
- M. Behrens, T. Lang, Extra-oral taste receptors—function, disease, and perspectives, *Front. Nutr.* 9 (2022) 881177, <https://doi.org/10.3389/fnut.2022.881177>.
- M. Behrens, W. Meyerhof, Oral and extraoral bitter taste receptors, in: W. Meyerhof, U. Beisiegel, H.-G. Joost (Eds.), *Sensory and Metabolic Control of Energy Balance*, 52, Springer Berlin Heidelberg, 2011, pp. 87–99, https://doi.org/10.1007/978-3-642-14426-4_8.
- A. Di Pizio, M.Y. Niv, Promiscuity and selectivity of bitter molecules and their receptors, *Bioorg. Med. Chem.* 23 (14) (2015) 4082–4091, <https://doi.org/10.1016/j.bmc.2015.04.025>.
- A. Di Pizio, Y. Ben Shoshan-Galeczki, J.E. Hayes, M.Y. Niv, Bitter and sweet tasting molecules: it's complicated, *Neurosci. Lett.* 700 (2019) 56–63, <https://doi.org/10.1016/j.neulet.2018.04.027>.
- J.A. Mennella, A.C. Spector, D.R. Reed, S.E. Coldwell, The bad taste of medicines: overview of basic research on bitter taste, *Clin. Ther.* 35 (8) (2013) 1225–1246, <https://doi.org/10.1016/j.clinthera.2013.06.007>.
- F.A. Shaik, N. Singh, M. Arakawa, K. Duan, R.P. Bhullar, P. Chelikani, Bitter taste receptors: extraoral roles in pathophysiology, *Int. J. Biochem. Cell Biol.* 77 (2016) 197–204, <https://doi.org/10.1016/j.biocel.2016.03.011>.
- D. Sun-Waterhouse, S.S. Wadhwa, Industry-relevant approaches for minimising the bitterness of bioactive compounds in functional foods: a review, *Food Bioprocess Technol.* 6 (3) (2013) 607–627, <https://doi.org/10.1007/s11947-012-0829-2>.
- A. Dagan-Wiener, A. Di Pizio, I. Nissim, M.S. Bahia, N. Dubovski, E. Margulis, M. Y. Niv, BitterDB: taste ligands and receptors database in 2019, *Nucleic Acids Res.* 47 (D1) (2019) D1179–D1185, <https://doi.org/10.1093/nar/ky974>.
- M. Malavolta, L. Pallante, B. Mavkov, F. Stojceski, G. Grasso, A. Korfiati, S. Mavroudi, A. Kalogeras, C. Alexakos, V. Martos, D. Amoroso, G. Di Benedetto, D. Piga, K. Theofilatos, M.A. Deriu, A survey on computational taste predictors, *Eur. Food Res. Technol.* 248 (9) (2022) 2215–2235, <https://doi.org/10.1007/s00217-022-04044-5>.
- E.-S.M. El-Kenawy, N. Khodadadi, S. Mirjalili, T. Makarovskikh, M. Abotaleb, F. K. Karim, H.K. Alkahtani, A.A. Abdelhamid, M.M. Eid, T. Horiuchi, E. Ibrahim, D. S. Khafaga, Metaheuristic optimization for improving weed detection in wheat images captured by drones, *Mathematics* 10 (23) (2022) 23, <https://doi.org/10.3390/math10234421>.
- E.-S.M. El-kenawy, N. Khodadadi, S. Mirjalili, A.A. Abdelhamid, M.M. Eid, A. Ibrahim, Greylag goose optimization: nature-inspired optimization algorithm, *Expert Syst. Appl.* 238 (2024) 122147, <https://doi.org/10.1016/j.eswa.2023.122147>.
- K.S. Gaber, M.K. Singla, Predictive analysis of groundwater resources using random forest regression, *Journal of Artificial Intelligence and Metaheuristics* (1) (2025) 11–19, <https://doi.org/10.54216/JAIM.090102>.
- M.E. Mohamed, A review on waste management techniques for sustainable energy production, *Metaheuristic Optimization Review* (2) (2025) 47–58, <https://doi.org/10.54216/MOR.030205>.
- N. Khodadadi, L. Abualigah, E.-S.M. El-Kenawy, V. Snaes, S. Mirjalili, An archive-based multi-objective arithmetic optimization algorithm for solving industrial engineering problems, *IEEE Access* 10 (2022) 106673–106698, <https://doi.org/10.1109/ACCESS.2022.3212081>.
- S. Zheng, M. Jiang, C. Zhao, R. Zhu, Z. Hu, Y. Xu, F. Lin, e-Bitter: bitterant prediction by the consensus voting from the machine-learning methods, *Front. Chem.* 6 (2018) 82, <https://doi.org/10.3389/fchem.2018.00082>.
- E. Margulis, A. Dagan-Wiener, R.S. Ives, S. Jaffari, K. Siems, M.Y. Niv, Intense bitterness of molecules: machine learning for expediting drug discovery, *Comput. Struct. Biotechnol. J.* 19 (2021) 568–576, <https://doi.org/10.1016/j.csbj.2020.12.030>.
- W. Bo, D. Qin, X. Zheng, Y. Wang, B. Ding, Y. Li, G. Liang, Prediction of bitterant and sweetener using structure-taste relationship models based on an artificial neural network, *Food Res. Int.* 153 (2022) 110974, <https://doi.org/10.1016/j.foodres.2022.110974>.
- G. Maroni, L. Pallante, G. Di Benedetto, M.A. Deriu, D. Piga, G. Grasso, Informed classification of sweeteners/bitterants compounds via explainable machine learning, *Curr. Res. Food Sci.* 5 (2022) 2270–2280, <https://doi.org/10.1016/j.crf.2022.11.014>.
- W. Huang, Q. Shen, X. Su, M. Ji, X. Liu, Y. Chen, S. Lu, H. Zhuang, J. Zhang, BitterX: a tool for understanding bitter taste in humans, *Sci. Rep.* 6 (2016) 1–8, <https://doi.org/10.1038/srep23450>.
- R. Tuwani, S. Wadhwa, G. Bagler, BitterSweet: building machine learning models for predicting the bitter and sweet taste of small molecules, *Sci. Rep.* 9 (1) (2019) 7155, <https://doi.org/10.1038/s41598-019-43664-y>.
- E. Margulis, Y. Slavutsky, T. Lang, M. Behrens, Y. Benjamini, M.Y. Niv, BitterMatch: recommendation systems for matching molecules with bitter taste receptors, *J. Cheminf.* 14 (1) (2022) 45, <https://doi.org/10.1186/s13321-022-00612-9>.
- F. Ferri, M. Cannariato, M.A. Deriu, L. Pallante, Machine learning approaches to predict TAS2R receptors for bitterants, *Biotechnol. Bioeng.* (2024) bit.28709, <https://doi.org/10.1002/bit.28709>.
- Z. Guo, K. Guo, B. Nan, Y. Tian, R.G. Iyer, Y. Ma, O. Wiest, X. Zhang, W. Wang, C. Zhang, N.V. Chawla, *Graph-based molecular representation learning* (No. arXiv: 2207.04869), arXiv (2023). <http://arxiv.org/abs/2207.04869>.
- S.M. Lundberg, G. Erion, H. Chen, A. DeGrave, J.M. Prutkin, B. Nair, R. Katz, J. Himmelfarb, N. Bansal, S.-I. Lee, *Explainable AI for trees: from local explanations to global understanding* (No. arXiv:1905.04610), arXiv (2019). <http://arxiv.org/abs/1905.04610>.
- P.E. Pope, S. Kolouri, M. Rostami, C.E. Martin, H. Hoffmann, Explainability methods for graph convolutional neural networks, in: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2019, pp. 10764–10773, <https://doi.org/10.1109/CVPR.2019.01103>.
- R. Ying, D. Bourgeois, J. You, M. Zitnik, J. Leskovec, *GNNExplainer: generating explanations for graph neural networks* (No. arXiv:1903.03894), arXiv (2019). <http://arxiv.org/abs/1903.03894>.
- S.M. Lundberg, S.-I. Lee, *A unified approach to interpreting model predictions*, in: I. Guyon, U.V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, R. Garnett (Eds.), *Advances in Neural Information Processing Systems*, 30, Curran Associates, Inc, 2017.

- [32] R.R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, D. Batra, Grad-CAM: visual explanations from deep networks via gradient-based localization, *Int. J. Comput. Vis.* 128 (2) (2020) 336–359, <https://doi.org/10.1007/s11263-019-01228-7>.
- [33] M. Behrens, M. Gu, S. Fan, C. Huang, W. Meyerhof, Bitter substances from plants used in traditional Chinese medicine exert biased activation of human bitter taste receptors, *Chem. Biol. Drug Des.* 91 (2) (2018) 422–433, <https://doi.org/10.1111/cbdd.13089>.
- [34] M. Cui, B. Chen, K. Xu, A. Rigakou, P. Diamantakos, E. Melliou, D.E. Logothetis, P. Magiatis, Activation of specific bitter taste receptors by olive oil phenolics and secoiridoids, *Sci. Rep.* 11 (1) (2021) 22340, <https://doi.org/10.1038/s41598-021-01752-y>.
- [35] T. Delompré, C. Belloir, C. Martin, C. Salles, L. Briand, Detection of bitterness in vitamins is mediated by the activation of bitter taste receptors, *Nutrients* 14 (19) (2022) 4141, <https://doi.org/10.3390/nu14194141>.
- [36] A. Jaggupilli, N. Singh, V.C. De Jesus, M.S. Gounni, P. Dhanaraj, P. Chelikani, Chemosensory bitter taste receptors (T2Rs) are activated by multiple antibiotics, *FASEB J.* 33 (1) (2019) 501–517, <https://doi.org/10.1096/fj.201800521RR>.
- [37] A. Karolkowski, C. Belloir, G. Lucchi, C. Martin, E. Bouzidi, L. Levavasseur, C. Salles, L. Briand, Activation of bitter taste receptors by saponins and alkaloids identified in faba beans (*Vicia faba* L. minor), *Food Chem.* 426 (2023) 136548, <https://doi.org/10.1016/j.foodchem.2023.136548>.
- [38] T. Lang, R. Lang, A. Di Pizio, V.K. Mittermeier, V. Schlagbauer, T. Hofmann, M. Behrens, Numerous compounds orchestrate coffee's bitterness, *J. Agric. Food Chem.* 68 (24) (2020) 6692–6700, <https://doi.org/10.1021/acs.jafc.0c01373>.
- [39] G. Morini, M. Winnig, T. Vennegeerts, G. Borgonovo, A. Bassoli, Vanillin activates human bitter taste receptors TAS2R14, TAS2R20, and TAS2R39, *Front. Nutr.* 8 (2021) 683627, <https://doi.org/10.3389/fnut.2021.683627>.
- [40] Fakhri Nouri, Sanadgol El-Senduny, Farzaei Abd-ElGhani, Chen, On the neuroprotective effects of naringenin: pharmacological targets, signaling pathways, molecular mechanisms, and clinical perspective, *Biomolecules* 9 (11) (2019) 690, <https://doi.org/10.3390/biom9110690>.
- [41] S. Soares, M.S. Silva, I. García-Estevéz, P. Großmann, N. Brás, E. Brandão, N. Mateus, V. De Freitas, M. Behrens, W. Meyerhof, Human bitter taste receptors are activated by different classes of polyphenols, *J. Agric. Food Chem.* 66 (33) (2018) 8814–8823, <https://doi.org/10.1021/acs.jafc.8b03569>.
- [42] S. Kim, J. Chen, T. Cheng, A. Gindulyte, J. He, S. He, Q. Li, B.A. Shoemaker, P. A. Thiessen, B. Yu, L. Zaslavsky, J. Zhang, E.E. Bolton, PubChem 2023 update, *Nucleic Acids Res.* 51 (D1) (2023) D1373–D1380, <https://doi.org/10.1093/nar/gkac956>.
- [43] A.P. Bento, A. Hersey, E. Félix, G. Landrum, A. Gaulton, F. Atkinson, L.J. Bellis, M. De Veij, A.R. Leach, An open source chemical structure curation pipeline using RDKit, *J. Cheminf.* 12 (1) (2020) 1, <https://doi.org/10.1186/s13321-020-00456-1>.
- [44] L. Pallante, A. Korfiati, L. Androustos, F. Stojceski, A. Bompotas, I. Giannikos, C. Raptopoulos, M. Malavolta, G. Grasso, S. Mavroudi, A. Kalogeras, V. Martos, D. Amoroso, D. Piga, K. Theofilatos, M.A. Deriu, Toward a general and interpretable umami taste predictor using a multi-objective machine learning approach, *Sci. Rep.* 12 (1) (2022) 21735, <https://doi.org/10.1038/s41598-022-25935-3>.
- [45] H. Moriwaki, Y.S. Tian, N. Kawashita, T. Takagi, Mordred: a molecular descriptor calculator, *J. Cheminf.* 10 (1) (2018) 4, <https://doi.org/10.1186/s13321-018-0258-y>.
- [46] S.G.K. Patro, K.K. Sahu, *Normalization: a preprocessing stage* (No. arXiv: 1503.06462), arXiv (2015), <https://doi.org/10.48550/arXiv.1503.06462>.
- [47] A. Natekin, A. Knoll, Gradient boosting machines, a tutorial, *Front. Neurobot.* 7 (2013). <https://www.frontiersin.org/articles/10.3389/fnbot.2013.00021>.
- [48] A.V. Dorogush, V. Ershov, A. Gulin, CatBoost: gradient boosting with categorical features support (No. arXiv:1810.11363), arXiv (2018), <https://doi.org/10.48550/arXiv.1810.11363>.
- [49] D.J. Rogers, T.T. Tanimoto, A computer program for classifying plants: the computer is programmed to simulate the taxonomic process of comparing each case with every other case, *Science* 132 (3434) (1960) 1115–1118, <https://doi.org/10.1126/science.132.3434.1115>.
- [50] Y. Akhlat, Y. Manzali, M. Chahhou, A. Zinedine, A new noisy random forest based method for feature selection, *Cybern. Inf. Technol.* 21 (2) (2021) 10–28.
- [51] F.J. Ferri, P. Pudil, M. Hatef, J. Kittler, Comparative study of techniques for large-scale feature selection, in: E.S. Gelsema, L.S. Kanal (Eds.), *Machine Intelligence and Pattern Recognition*, 16, North-Holland, 1994, pp. 403–413, <https://doi.org/10.1016/B978-0-444-81892-8.50040-7>. * *This work was supported by a SERC grant GR/E 97549. The first author was also supported by a FPI grant from the Spanish MEC, PF92 73546684.
- [52] T.N. Kipf, M. Welling, Semi-supervised classification with graph convolutional networks (No. arXiv:1609.02907; version 4), arXiv (2017), <https://doi.org/10.48550/arXiv.1609.02907>.
- [53] M. Fey, J.E. Lenssen, Fast graph representation learning with PyTorch geometric. <https://doi.org/10.48550/ARXIV.1903.02428>, 2019.
- [54] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, S. Chintala, *PyTorch: an imperative style, high-performance deep learning library* (No. arXiv:1912.01703), arXiv (2019). <http://arxiv.org/abs/1912.01703>.
- [55] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, P.S. Yu, A Comprehensive survey on graph neural networks, *IEEE Transact. Neural Networks Learn. Syst.* 32 (1) (2021) 4–24, <https://doi.org/10.1109/TNNLS.2020.2978386>.
- [56] A.A. Hagberg, D.A. Schult, P.J. Swart, Exploring network structure, dynamics, and function using NetworkX, in: G. Varoquaux, T. Vaught, J. Millman (Eds.), *Proceedings of the 7th Python in Science Conference, 2008*, pp. 11–15.
- [57] Y. He, K. Liu, Y. Liu, W. Han, Prediction of bitterness based on modular designed graph neural network, *Bioinformatics Advances* 4 (1) (2024) vbae041, <https://doi.org/10.1093/bioadv/vbae041>.
- [58] J. Lim, S. Ryu, K. Park, Y.J. Choe, J. Ham, W.Y. Kim, Predicting drug–target interaction using a novel graph neural network with 3D structure-embedded graph representation, *J. Chem. Inf. Model.* 59 (9) (2019) 3981–3988, <https://doi.org/10.1021/acs.jcim.9b00387>.
- [59] T. Nguyen, H. Le, T.P. Quinn, T. Nguyen, T.D. Le, S. Venkatesh, GraphDTA: predicting drug–target binding affinity with graph neural networks, *Bioinformatics* 37 (8) (2021) 1140–1147, <https://doi.org/10.1093/bioinformatics/btaa921>.
- [60] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, Y. Bengio, Graph attention networks (No. arXiv:1710.10903), arXiv (2018), <https://doi.org/10.48550/arXiv.1710.10903>.
- [61] C. Agarwal, O. Queen, H. Lakkaraju, M. Zitnik, Evaluating explainability for graph neural networks. <https://doi.org/10.48550/ARXIV.2208.09339>, 2022.
- [62] S. Zheng, W. Chang, W. Xu, Y. Xu, F. Lin, e-Sweet: a machine-learning based platform for the prediction of sweetener and its relative sweetness, *Front. Chem.* 7 (JAN) (2019) 1–14, <https://doi.org/10.3389/fchem.2019.00035>.
- [63] W. Xu, L. Wu, S. Liu, X. Liu, X. Cao, C. Zhou, J. Zhang, Y. Fu, Y. Guo, Y. Wu, Q. Tan, L. Wang, J. Liu, L. Jiang, Z. Fan, Y. Pei, J. Yu, J. Cheng, S. Zhao, T. Hua, Structural basis for strychnine activation of human bitter taste receptor TAS2R46, *Science* 377 (6612) (2022) 1298–1304, <https://doi.org/10.1126/science.abo1633>.