Using Molecular Embeddings in QSAR Modeling: Does it Make a Difference?

María Virginia Sabando,*,†,‡ Ignacio Ponzoni,†,‡ Evangelos E. Milios,¶ and Axel J. Soto†,‡

†*Institute for Computer Science and Engineering (UNS-CONICET), Bahía Blanca, Argentina*

‡*Department of Computer Science and Engineering, Universidad Nacional del Sur, Bahía Blanca, Argentina*

¶*Faculty of Computer Science, Dalhousie University, Halifax, NS, Canada*

E-mail: virginia.sabando@cs.uns.edu.ar

Abstract

With the consolidation of deep learning in drug discovery, several novel algorithms for learning molecular representations have been proposed. Despite the interest of the community in developing new methods for learning molecular embeddings and their theoretical benefits, comparing molecular embeddings with each other and with traditional representations is not straightforward, which in turn hinders the process of choosing a suitable representation for QSAR modeling. A reason behind this issue is the difficulty of conducting a fair and thorough comparison of the different existing embedding approaches, which requires numerous experiments on various datasets and training scenarios. To close this gap, we reviewed the literature on methods for molecular embeddings and reproduced three unsupervised and two supervised molecular embedding techniques recently proposed in the literature. We compared these five

1

methods concerning their performance in QSAR scenarios using different classification and regression datasets. We also compared these representations to traditional molecular representations, namely molecular descriptors and fingerprints. As opposed to the expected outcome, our experimental setup consisting of over $25,000$ trained models and statistical tests revealed that the predictive performance using molecular embeddings did not significantly surpass that of traditional representations. While supervised embeddings yielded competitive results compared to those using traditional molecular representations, unsupervised embeddings tended to perform worse than traditional representations. Our results highlight the need for conducting a careful comparison and analysis of the different embedding techniques prior to using them in drug design tasks, and motivate a discussion about the potential of molecular embeddings in computer-aided drug design.

Introduction

Quantitative Structure-Activity Relationship (QSAR) models constitute the cornerstone of modern in silico drug discovery.[1–3] QSAR models are regression or classification models that predict the relationship between molecular features of compounds and a target property. They are extensively used in modern drug discovery to accelerate drug candidate identification while reducing costs associated with molecular synthesis and wet lab experiments. QSAR models are usually trained using traditional molecular representations, such as molecular descriptors and fingerprints, which are computed using widely known algorithms.[4] Although these molecular representations are well-established and often yield good classification results, there is an increasing interest in exploring new, enriched representations that take into account multiple aspects of the molecule.[5]

Deep learning techniques applied to Natural Language Processing (NLP), originally developed for text, have experienced a prolific development during the last decade and are being increasingly applied to other domains.[6–8] In drug design, significant research efforts

have been put on designing novel techniques for learning rich molecular representations during the last years.[5,6,9] These representations, which we refer to as *molecular embeddings*, can be trained in different ways to capture diverse information about physical-chemical properties, molecular structure and bioactivity. In addition, unsupervised representation learning methods inspired by NLP approaches can be adapted to use huge amounts of molecular information in SMILES format,[10] thus benefiting from the inherent diversity in large ensembles of unlabeled compounds.

Algorithms for learning molecular embeddings employ a wide variety of state-of-the-art deep learning techniques,[6] ranging from neural-based autoencoders to graph neural networks[8,11] and self-attention.[12] Molecular embedding algorithms vary in complexity and intricacy, allowing to tailor the learned representations to specific tasks. For instance, many recent self-attention methods for molecular embeddings are designed to identify the molecular substructures that have a significant impact on their bioactivity profile.[13,14] Nowadays, several QSAR studies employ molecular embeddings instead of traditional representations as training data.[5,15,16] As depicted in Figure 1, molecular embeddings are usually fixed-size dense vectors of real numbers, whereas traditional molecular representations can either consist of real number vectors or bit vectors varying in sparsity. These representations are fed to classification or regression models for QSAR modeling.

At learning time, molecular embeddings can either be *supervised*, meaning that they take into account labeled/external information about the bioactivity profile of the molecules during training, or *unsupervised*, which means that the embeddings are built with no human-labeled information. While supervised embeddings encode bioactivity or physical-chemical information, which could favor them in predictive tasks, they are consequently less flexible than unsupervised embeddings, since new embeddings need to be learned for different biological targets. Moreover, in order to learn supervised embeddings it is necessary to have labeled datasets, which are often small and scarce, thus potentially affecting the embedding quality negatively.
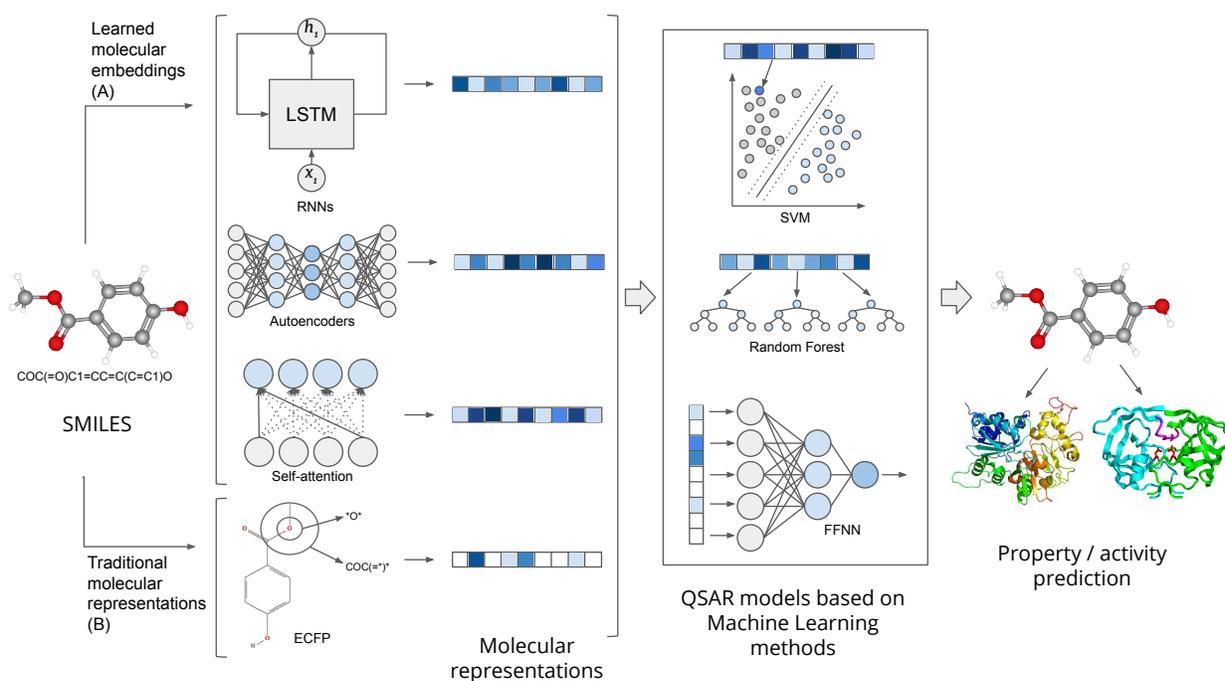
Figure 1: QSAR models can either be trained using learned molecular embeddings derived from deep learning and NLP techniques (A), or employing traditional molecular representations obtained from a *feature engineering* process (B). Molecular embeddings are dense vectors of real numbers, whereas traditional molecular representations can be either real number vectors or bit vectors varying in sparsity. QSAR models are classification or regression models that are trained with these representations.

Despite the proliferation of learned molecular representation methods, there is no empirical research on how to choose the most suitable representation method for QSAR analysis. Interestingly, some recent studies have been published where molecular embeddings appear to only match or slightly surpass traditional representations in QSAR modeling.[16–22] While establishing a fair comparison between molecular embeddings and traditional representations is not straightforward, we argue that such comparison is necessary and that it needs to be carried out through an extensive and careful experimental workflow. In order to establish a comparison on a fair ground, we propose a carefully designed evaluation where hyperparameters are optimized thoroughly, datasets of different characteristics are used, and several models are trained to account for their inherent stochasticity. In particular, this work aims to address the following research questions:

- Q1: What are the main molecular embedding methods used for QSAR modeling in the literature? Do they outperform traditional molecular representations in a classification/regression task?

- Q2: Does incorporating information about the biological target into the molecular embedding (supervised embeddings) incur higher predictive performance than that obtained from unsupervised embeddings?

- Q3: Do different preprocessing decisions, such as the canonical form of the SMILES formulas used as training data or the size of the final embeddings, significantly impact on the predictive performance of a QSAR model using molecular embeddings?

In this scenario, we conducted an extensive analysis of different molecular embedding techniques based on state-of-the-art deep learning methods focusing on their suitability for QSAR modeling. We reproduced five recently proposed embedding techniques: two *supervised* and three *unsupervised* techniques. All five techniques were trained using SMILES formulas as the raw molecular representation. To train the unsupervised methods, we employed a dataset of 40 million randomly selected and curated compounds retrieved from the

ZINC database.[23] In addition, we tested two different ways of computing canonical SMILES formulas in an effort to leverage the way in which each unsupervised technique processes the SMILES formulas. We trained different classifiers using eight labeled datasets with varying class imbalance levels and sizes, corresponding to both classification and regression tasks. Finally, we contrasted the performance attained by molecular embeddings to the baseline results obtained using traditional molecular representations. Our experimental workflow consisted of over 25,000 trained models, which comprised several stages of model selection and replications, followed by statistical analyses of the results.

Related work

The trustworthiness of QSAR models and their predictive performance is linked to the choice of the molecular representation used for training.[5,24] For this reason, for many years the standard has been to manually engineer high-quality molecular representations. This process, known as *feature engineering*, has led to widespread and commonly used molecular representations.[4,25] Traditional molecular representations vary mostly in the type of information they encode and their use depends on the specific task.[26,27] Among the most commonly used traditional molecular representations, we can enumerate: *molecular descriptors*,[4] *Extended Connectivity Fingerprints (ECFPs)*[28] and *Molecular ACCess System (MACCS)* keys.[29]

While a vast number of published QSAR models are based on these traditional representations,[30–33] the process of obtaining such representations via feature engineering is costly and requires strong domain-specific expertise. Moreover, since each of them encodes different information, there is no single representation suitable for every task. For this reason, there is an increasing trend towards the use of versatile molecular representations that capture diverse aspects of the chemical space,[5] so many new methods for learning molecular embeddings have been developed recently, where most of them are based on deep learning techniques.[5,6]

Recent methods for learning molecular embeddings employ *SMILES* formulas,[10] which is the most widely used linear representation for encoding molecular graph information. Since SMILES formulas directly encode the molecular graph into a sequence of ASCII characters, they can be used by deep learning techniques designed for sequential or graph-like data.[6,16] Given their high availability—most molecular databases are stored in SMILES format—, many unsupervised methods for learning molecular representations have been developed and trained on large databases without the need for labeled information.[17,34–37] While images or structured graph representations have also been employed for learning molecular embeddings in QSAR modeling,[19,38,39] their usage is not as widespread and has limitations in terms of data availability. Thus, in this paper, we focused on SMILES-based embedding methods.

The high availability of molecular data in SMILES format has motivated numerous approaches for learning molecular embeddings based on autoencoders, which enable to conduct unsupervised training procedures.[17,18,37,40] In particular, Öztürk et al.[36] introduced *SMILESVec*, an unsupervised SMILES-based method that learns representations for small molecules using the popular *word2vec* model by Mikolov et al.[41] The authors performed a *tokenization* step onto a dataset of SMILES formulas—i.e., converted each SMILES into a string of distinguishable tokens in an alphabet— by computing overlapping SMILES substrings, which were then used to train their model. Finally, they computed an embedding for each molecule in the corpus by averaging the learned vectors for each of the tokens in its SMILES formula. Jaeger et al.[17] presented *Mol2Vec*, which is also based on *word2vec*.[41] In this method, SMILES formulas undergo a preprocessing stage before being fed to the autoencoder, which consists of tokenizing the SMILES formulas using the Morgan algorithm for ECFPs.[28] After an unsupervised training phase, the final representations are computed by summing the embedded vectors of all tokens in the molecule.

Since SMILES is a linear molecular notation of a sequential nature, methods for either generating new molecules or learning new molecular embeddings based on Recurrent Neural Networks (RNNs) are also found in the literature.[37,42,43] Most generative models in the

literature are based on RNN autoencoders trained on large sets of unlabeled data,[6] which constitutes an interesting approach for learning molecular representations. In this direction, Xu et al.[37] proposed *Seq2Seq Fingerprint*, an unsupervised method based on a multi-layer RNN autoencoder built using Gated Recurrent Unit cells (GRUs).[44] *Seq2Seq Fingerprint* learns directly from SMILES formulas leveraging information from long-term dependencies and sequential relationships between characters. The molecular embeddings are finally obtained by concatenating the hidden states of the autoencoder.

A recent trend in learned molecular representations is the use of self-attention,[12] which enables to capture rich information about relevant substructures of the molecule, while avoiding the lengthy training procedures of RNNs.[12] Self-attention strongly relates to the way GNNs operate on graph-like data[45] but are designed to learn from sequence data such as SMILES formulas. In this direction, Oskooei et al.[13] developed *PaccMann*, a novel approach for predicting anticancer compound sensitivity by means of multi-modal attention-based neural networks. The authors proposed a series of supervised encoders; the best performing one was a self-attentive encoder trained using a simple tokenization of SMILES strings and complemented with gene expression information. A shallow feed-forward neural network for property prediction was stacked to the embedding model so that the molecular embeddings were learned during the training phase of the classifier or regression model. Zheng et al.[14] developed *SA-BiLSTM*, a supervised method based on a combination of a Bidirectional Long-Short Term Memory (LSTM)[46] RNN and a self-attention mechanism. This model is then plugged into a shallow feed-forward neural network for property prediction, thus the self-attentive embeddings are learned along with the prediction task. One interesting aspect of this method is that it does not learn from SMILES formulas directly but instead uses sequences of token embeddings from a pretrained *Mol2Vec* model.[17] Wang et al.[47] developed *SMILES-BERT*, a semi-supervised model consisting of a Transformer Layer[12] trained using a large unlabeled dataset through a masked SMILES recovery task, as in the popular NLP *BERT* architecture.[48] This model can be then fine-tuned by means of a labeled dataset.

8

Although *SMILES-BERT* seems a promising approach, important implementation details are missing to fully reproduce the method, especially in terms of the fine-tuning stage.

Materials and methods

In order to enable reproducibility of this paper, in this section we provide a thorough explanation of our experimental setup and a workflow overview, including a description of the datasets and their preprocessing stage, the neural architectures of the five methods we reproduced and their training. Finally, we provide details on the classification and regression stages by which we evaluated the learned molecular embeddings. Further details about materials and source code can be found in Section "Data availability statement".

Datasets

Following the steps needed to reproduce the unsupervised techniques, we collected and downloaded the SMILES formulas of approximately 200 million purchasable compounds from ZINC database.[23] We conducted a preprocessing stage that consisted of filtering out compounds that did not comply with Lipinski's Rule of 5.[49] We only kept compounds having a molecular weight between 12 and 600, heavy atom count between 3 and 50 and *cLogP* (partition coefficient octanol/water) value between -5 and 7. We also filtered out compounds that presented non-drug-like atoms, such as heavy metals, and removed salts and solvents. Finally, we obtained the canonical SMILES formulas for each of the remaining compounds. This whole process was carried out using RDKit.[50] After the preprocessing stage, we randomly selected a subset of 40 million compounds—twice as many compounds as the ones used in the reference papers[17,36,37]—, which were later used to train the unsupervised methods.

We also selected eight different labeled datasets, five classification datasets and three regression datasets, which were used to train the supervised methods and to evaluate all embedding methods. All of the classification datasets pose binary classification tasks. We

prioritized datasets that had initially been used in the reference papers[13,14,17,36,37] while we accounted for diversity of sizes and class imbalance scenarios. The selected datasets were preprocessed following the same procedure used for ZINC, described above. Further details about class imbalance and numbers of compounds are summarized in Table 1. Namely, these datasets are:

- *SR-ARE*, a bioassay for small molecule agonists of the antioxidant response element (ARE) signaling pathway[i]. This assay of data is contained in the *Tox21 challenge* dataset, consisting of qualitative toxicity measurements on twelve biological targets[ii].

- *SR-MMP*, a stress response assay for small molecule disruptors of the mitochondrial membrane potential (MMP)[iii]. This dataset is also included in the *Tox21 challenge* dataset.

- *SR-ATAD5*, a set of small molecules that induce genotoxicity in human embryonic kidney cells expressing luciferase-tagged ATAD5[iv]. This dataset also belongs to the *Tox21 challenge* dataset.

- *HIV*, a dataset introduced by the Drug Therapeutics Program (DTP) AIDS Antiviral Screen containing information about molecular ability to inhibit HIV replication. Screening results were categorized as confirmed inactive (CI), confirmed active (CA) and confirmed moderately active (CM). We merged the compounds categorized under the latter two labels, which yielded two classes: inactive (CI) and active (CA and CM)[v].

- *PCBA-686978*, a PubChem bioassay containing information about molecular ability to inhibit the human tyrosyl-DNA phosphodiesterase 1 (TDP1)[vi].

---

[i]https://pubchem.ncbi.nlm.nih.gov/bioassay/743219
[ii]https://tripod.nih.gov/tox21/challenge/about.jsp
[iii]https://pubchem.ncbi.nlm.nih.gov/bioassay/720637
[iv]https://pubchem.ncbi.nlm.nih.gov/bioassay/720516
[v]https://wiki.nci.nih.gov/display/NCIDTPdata/AIDS+Antiviral+Screen+Data
[vi]https://pubchem.ncbi.nlm.nih.gov/bioassay/686978

- *ESOL*, a dataset consisting of water solubility data for 1128 compounds, used to train models that estimate solubility directly from chemical structures.[51]

- *FreeSolv*—the *Free Solvation Database*—, a dataset comprising experimental and calculated hydration free energy values for small molecules in water.[52]

- *Lipophilicity*, a dataset curated from ChEMBL database[53] including experimental results of octanol/water distribution coefficient (*LogD 7.4*) for 4200 compounds [vii].

Table 1: Details of labeled datasets. The imbalance ratio for classification datasets is computed as the number of active compounds every 100 inactive compounds.

Dataset	Task	# Compounds	# Active	# Inactive	Imbalance ratio
SR-ARE	classification	5956	941	5015	18.76
SR-MMP	classification	5937	925	5012	3.68
SR-ATAD5	classification	7251	258	6993	18.46
HIV	classification	41127	1443	39684	27.50
PCBA-686978	classification	302175	62800	239375	3.81
ESOL	regression	1128	-	-	-
FreeSolv	regression	642	-	-	-
Lipophilicity	regression	4200	-	-	-

Molecular embedding methods

After identifying the most prominent deep learning techniques for learning molecular representations based on SMILES formulas, we chose *SMILESVec*,[36] *Mol2Vec*,[17] *Seq2Seq Fingerprint*[37]— which we refer to as *Seq2Seq* in this paper—, *PaccMann's* self attentive encoder[13]—named *PaccMann* throughout this paper— and *SA-BiLSTM* encoder[14] as our reference embedding methods. We decided not to experiment with *SMILES-BERT*, due to the reproducibility issues pointed out in the previous section and because both *SA-BiLSTM* and *PaccMann* are also self-attention-based encoders. We selected these methods to test different deep learning architectures, model sizes in terms of trainable parameters and considering

---

[vii]https://www.ebi.ac.uk/chembl/document_report_card/CHEMBL3301361/

both supervised and unsupervised training scenarios. Moreover, we aimed at conducting a training phase consisting of multiple runs of each method in order to account for their intrinsic variance. We conducted our benchmarking process by testing an extensive range of combinations of hyperparameters for each embedding method and classification/regression model and by leveraging different canonical forms of the SMILES formulas in the datasets. All of the unsupervised methods were trained on the 40 million compound dataset retrieved and curated from ZINC, whereas the supervised methods were trained using each of the eight labeled datasets described previously. We hereby describe the neural-based architectures of each embedding method and the main characteristics of the obtained molecular representations. We provide further details on the parameterization of each model in Appendices A and B.

Unsupervised methods

The architecture of *SMILESVec*[36] consists of a feed-forward neural-based autoencoder [viii], as proposed by Mikolov et al.[41] According to the original paper, we first tokenized the SMILES formulas into overlapping substrings of eight characters and compiled a vocabulary set with all the tokens in the training set. Considering the long-term dependencies often present in the SMILES sequences, such as branches or stereochemistry information, and that the eight-gram tokenization might limit the ability of the model to learn from them, we employed two different canonicalization processes on the ZINC training data, thus obtaining two different training sets for this method. On the one hand, we employed RDKit canonical SMILES,[50] whereas, on the other hand, we tested *DeepSMILES* canonical formulas,[54] which are built by eliminating cycles in the molecular graph in order to reduce long-term dependencies in the sequences. *DeepSMILES* has also been employed by the authors who proposed *SMILESVec* on a follow-up study,[40] which showed improved performance over the results obtained with traditional SMILES formulas.

[viii]`https://github.com/hkmztrk/SMILESVecProteinRepresentation/tree/master/source/word2vec`.

As proposed by the *SMILEVec* authors,[36] a Skip-gram *word2vec* model[41] was used, in which the size of the embeddings is determined by the number of nodes in its hidden layer. We trained two different embedding sizes for each canonicalization technique mentioned above: 100 and 300. While the embedding size used in the original paper was 100, we decided to test 300-dimensional embeddings to compare *SMILESVec* against higher-dimensional embeddings obtained by other methods tested in this paper. As performed by Öztürk et al.,[36] the unsupervised embedding models were trained for 20 epochs and the final molecular representations were obtained by computing the average of the learned vectors for each token in the SMILES formulas.

Mol2Vec[17] is also an unsupervised molecular embedding method based on *word2vec*.[41] The architecture of the model reported by the authors [ix] consists of a Skip-gram *word2vec* model, where the embedding size matches the size of the hidden layer of the autoencoder. Following the preprocessing steps indicated by Jaeger et al.,[17] we tokenized the RDKit canonical SMILES formulas in the ZINC database into sequences of chemical words obtained by employing an adaptation of the Morgan algorithm for computing ECFPs.[28] We did not use *DeepSMILES* canonical SMILES to train this method since the Morgan algorithm cannot reconstruct the molecular graph from them. Following the reference paper, we compiled a vocabulary set including all chemical words obtained from the ZINC training dataset and trained two models for 5 epochs each, with embedding sizes 100 and 300, respectively. Finally, we obtained the molecular embeddings by computing the sum of the learned vectors of all tokens in each molecule.

The third and last unsupervised method we reproduced is *Seq2Seq*,[37] an embedding model consisting of an RNN-based autoencoder[x]. The SMILES formulas in the ZINC training set were canonicalized to both RDKit and *DeepSMILES* canonical SMILES. While RDKit canonical SMILES were the only ones used in the reference paper,[37] we decided to additionally employ *DeepSMILES* canonical formulas, aiming to test whether they contributed to

---

[ix]https://github.com/samoturk/mol2vec.
[x]https://github.com/XericZephyr/seq2seq-fingerprint.

the learning process of the model. The sequences in both canonical datasets were tokenized by splitting the SMILES strings into separate characters and padded to match the length of the longest sequence in the dataset, as performed in the reference paper. According to Xu et al.,[37] the architecture of the model consists of a multi-layer bidirectional GRU autoencoder, and the learned embeddings are obtained by concatenating the hidden states of the trained model. Thus, their size is determined by the number of hidden units and the number of autoencoder layers. We trained two different models for each of the canonical forms of SMILES strings, obtaining 100-dimensional and 384-dimensional embeddings. Size 384 was the default embedding size reported by the authors of the method, and we also tested size 100 to account for a fair comparison with *SMILESVec* and *Mol2Vec*. We provide a summary of the molecular embeddings we obtained after training all the mentioned variants of the unsupervised methods in Table 2.

Table 2: Summary of all unsupervised embedding methods. The column *Denomination* denotes the name by which we refer to each of the learned embeddings throughout this paper. The denominations marked with an asterisk (*) correspond to the embeddings we decided to experiment with in addition to the ones proposed in the reference papers.

Method	Canonicalization	Embedding size	Denomination
SMILESVec	RDKit	100	SMILESVec_100
		300	SMILESVec_300 (*)
	DeepSMILES	100	Deep_SMILESVec_100
		300	Deep_SMILESVec_300 (*)
Mol2Vec	RDKit	100	Mol2Vec_100
		300	Mol2Vec_300
Seq2Seq	RDKit	100	Seq2Seq_100 (*)
		384	Seq2Seq_384
	DeepSMILES	100	Deep_Seq2seq_100 (*)
		384	Deep_Seq2seq_384 (*)

Supervised methods

PaccMann[13] is a supervised embedding method based on a set of multi-modal neural-based encoders that learn from both SMILES formulas and gene expression information, originally designed for predicting anticancer compound sensitivity. In particular, we reproduced the *self-attention (SA)* encoder proposed by the authors, which reportedly yielded the best results. Unlike the reference paper, we did not employ gene expression information in the training phase and instead trained the self-attentive encoder using only RDKit canonical SMILES formulas. In order to attain such a model, we adapted the source code provided by Oskooei et al.[13][xi]. According to the reference paper, the SMILES formulas were tokenized following an algorithm[55] that consists of separating the SMILES formulas into characters and performing a basic filtering step. After such tokenization, each sequence was padded to match the longest sequence in the dataset.

The architecture of *PaccMann* consists of an input layer by which the tokenized SMILES formulas are fed to the model. Next, a sinusoidal positional encoding function[12] is optionally applied to these inputs, followed by a single-headed self-attention layer, whose implementation details are in the reference paper.[13] The outcome of the self-attentive layer is fed to a shallow feed-forward neural network for property prediction. The learned molecular embeddings were extracted from the output of the self-attention layer, and their size is computed as the product between the number of hidden units u in the self-attention layer and the length of the input sequence n. Consequently, the embedding size $n \times u$ is different for each of the eight labeled datasets on which this method was trained. A summary of the embedding dimensionality for each labeled dataset is provided in Table 3.

The second supervised method we reproduced in this paper is *SA-BiLSTM*.[14] This method consists of a self-attentive model combined with a bidirectional RNN using LSTM nodes and followed by a shallow feed-forward neural network for property prediction. According to the reference paper, we trained this model using *Mol2Vec* embeddings[17] as input

---

[xi]https://github.com/drugilsberg/paccmann.

data, which were computed for the eight labeled datasets described previously and using a pretrained model provided by Jaeger et al.[17] We employed RDKit canonical SMILES to train this method. Unlike the other methods tested in this paper, the authors of *SA-BiLSTM* have not provided the source code or an installable package. Therefore, we implemented the model from the equations and details provided in their paper.

According to the authors, the architecture of *SA-BiLSTM* consists of an input layer that receives the *Mol2Vec* embeddings, followed by a single-layer bidirectional LSTM RNN. The output of such a layer is then passed to a multi-head self-attention layer, which yields the final self-attentive embeddings of the compounds. These embeddings are fed to a shallow feed-forward neural network for property prediction, consisting of a single fully-connected layer. The self-attentive embeddings obtained from this model consist of a set of r vectors of size $2u$, where r is the number of attention heads and u is the number of hidden units in the LSTM layer (the factor of 2 is because the RNN is bidirectional). According to the reference paper, these embeddings (the r vectors) can be analyzed separately. However, since we wanted to use the learned information from all the attention heads for prediction, we concatenated them for the subsequent classification/regression step. Thus, the resulting embeddings consist of vectors of size $r \times 2u$. Similarly to the case of *PaccMann*, because of the fitting process the size of the embedding varies for each of the eight labeled datasets used to train *SA-BiLSTM*. A summary of the *SA-BiLSTM* embedding dimensionality for each labeled dataset is provided in Table 3.

Experimental design

In this section, we give an overview of the experimental workflow we conducted to obtain and evaluate the molecular embeddings computed from each of the reviewed methods. A summary of such workflow can be found in Figure 2.
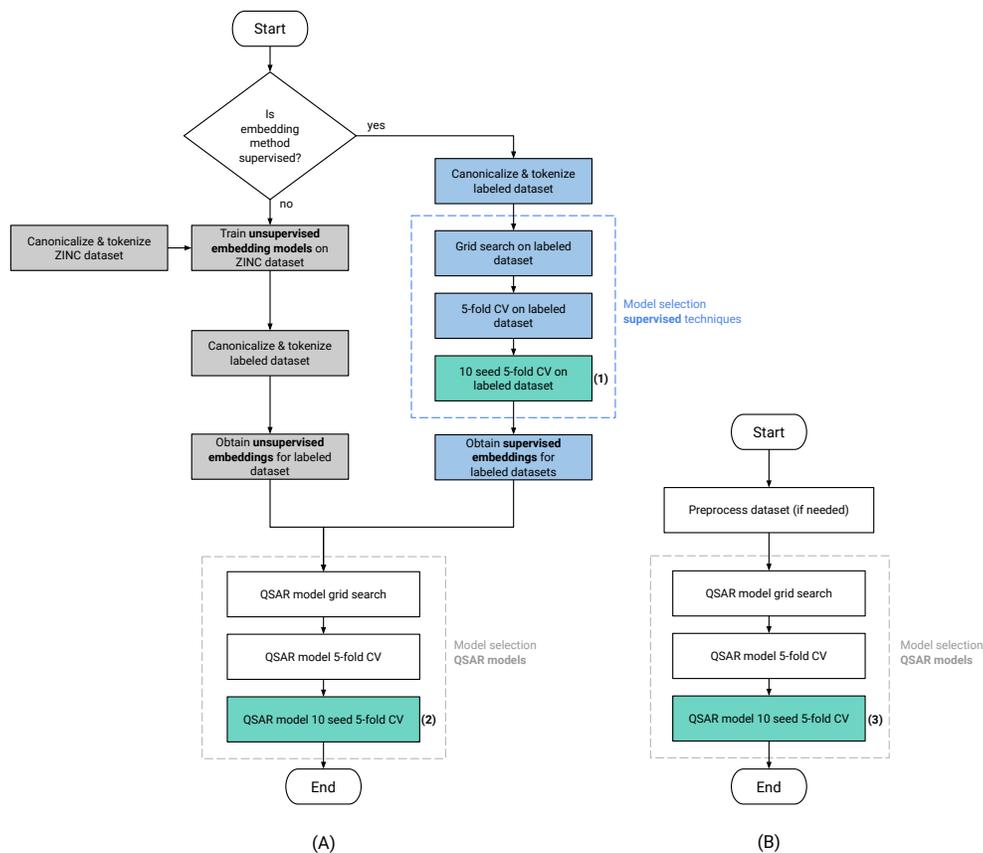
Figure 2: Overview of the experimental workflow: (A) corresponds to the workflow for embedding-based representations, whereas (B) depicts the workflow for traditional molecular representations. The diagram cells in grey correspond to the experiments on the unsupervised embedding techniques, whereas the diagram cells in blue correspond to the experiments on the supervised embedding techniques. The three final cells correspond to the prediction stage. The cells in green indicate the reported classification/regression results: (1) the *fitting* results, i.e., the classification/regression results obtained from the supervised embedding methods as a result of their training process (2) the classification/regression results obtained from the QSAR models trained using either supervised or unsupervised embeddings. (3) the classification/regression results obtained from the QSAR models trained using traditional molecular representations.

Table 3: Summary of all supervised embedding methods and their embedding sizes.

Method/Denomination	Dataset	Embedding size
PaccMann	SR-ARE	12000
	SR-MMP	24000
	SR-ATAD5	24000
	HIV	24200
	PCBA-686978	6976
	ESOL	9700
	FreeSolv	5300
	Lipophilicity	10250
SA-BiLSTM	SR-ARE	2560
	SR-MMP	1920
	SR-ATAD5	1280
	HIV	2560
	PCBA-686978	1280
	ESOL	1280
	FreeSolv	1280
	Lipophilicity	1280

Training and embedding extraction

The first step of our experimental setup consisted of training a model for each of the reviewed embedding methods. The training phase of the unsupervised methods (*SMILESVec, Mol2Vec* and *Seq2Seq)* was simpler than that of the supervised methods since the training dataset is unlabeled, and thus no early stopping criteria to avoid overfitting is needed. We trained each method following the amount of time or epochs and hyperparameterization specified in each reference paper. We obtained ten unsupervised embedding models for each combination of reviewed unsupervised method, embedding size and SMILES canonicalization of the ZINC training dataset, as summarized in Table 2. Finally, we tokenized the SMILES formulas in each labeled dataset as required by each embedding approach, and extracted the molecular embeddings from the ten embedding models.

Regarding the supervised methods (*PaccMann* and *SA-BiLSTM*), we conducted a broader range of experiments during their training stage compared to the experiments on the unsupervised methods. Our goal was to attain embedding models whose performance could

not be attributed solely to variance in the data partitions, in the initialization of the weights of the model or to nuances in the chosen hyperparameters. For this reason, we designed the following training workflow:

1. We filtered and tokenized the compounds of the five labeled datasets according to the requirements of each of the supervised methods.

2. We performed a model selection stage consisting of two steps:

   (a) We conducted a hyperparameter *grid search*: an exploratory search for the best performing hyperparameter combination for each method. The ranges of hyperparameter values included those tested in the reference papers[13,14] plus other values that might improve the training results. We selected the combinations of hyperparameters that yielded top results on each labeled dataset. A complete list of all hyperparameters tested on *PaccMann* and *SA-BiLSTM* is provided in Appendix A.

   (b) Each of the hyperparameter combinations selected in the previous step was further tested on a stratified five-fold cross-validation process, in order to confirm that the results found during the grid search were not imputable only to the data partition used during that process. As a result, we selected a single combination of hyperparameters for each labeled dataset based on the average validation results of the five-fold cross-validation process.

3. We performed a *replication stage*: we trained ten trials of five-fold cross-validation of the model selected during the model selection stage. Each trial used a different random seed for weights initialization. This step was carried out to rule out any artifacts in the results due to the random weight initialization of the embedding models in the previous steps. As a result of this step, we obtained the classification/regression results obtained from the inherent training process of the supervised embedding methods, which we refer to as *fitting* results.

4. Since a single embedding model was needed to extract the supervised embeddings, we trained an embedding model using a stratified data partition (80% train – 20% internal validation for early stopping). Finally, we extracted the molecular representations for each labeled dataset from these embedding models.

In addition to Figure 2, which summarizes the whole experimental workflow of our paper, we provide a more detailed graphical summary of the training stage of both unsupervised and supervised methods in Figures 3 and 4, respectively. The source code needed to extract the embeddings from the trained embedding models is either provided by the original authors, as in the case of *Seq2Seq*[xii], or by us. Please refer to Section "Data availability statement" for further details.

Evaluation of the molecular embeddings

In order to evaluate the different learned molecular embeddings in QSAR modeling, we tested them on five classification and three regression tasks, defined by the eight labeled datasets described in Table 1. We obtained a total of ten different unsupervised molecular embeddings and two different supervised molecular embeddings per labeled dataset, as shown in Figures 3 and 4. In addition, we computed three traditional molecular representations for each labeled dataset: a 1024-bit ECFP4 fingerprint using RDKit,[50] a 166-bit MACCS keys fingerprint also computed with RDKit, and a vector of molecular descriptors computed using Mordred.[56] We computed 0D, 1D and 2D descriptors for all datasets and discarded those having more than 5% of *NaN* entries. The final number of molecular descriptors is as follows: 1018 for SR-ARE, 1016 for SR-MMP and SR-ATAD5, 1150 for HIV, 1428 for PCBA-686978, 1226 for ESOL, 1176 for FreeSolv and 1429 for Lipophilicity. All the traditional representations were computed from RDKit canonical SMILES formulas.

In summary, we computed a total of 15 (10 + 2 + 3) representations for each dataset, which we then used to train our QSAR models. For each of the classification tasks we trained

---

[xii]https://github.com/XericZephyr/seq2seq-fingerprint.
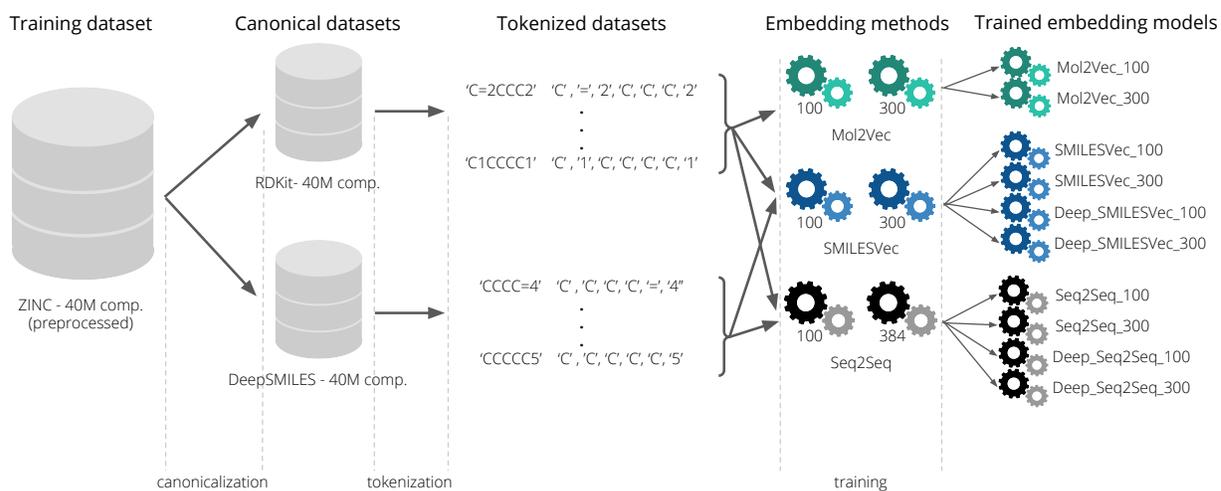
(a) Unsupervised training using 40 million compounds from ZINC dataset
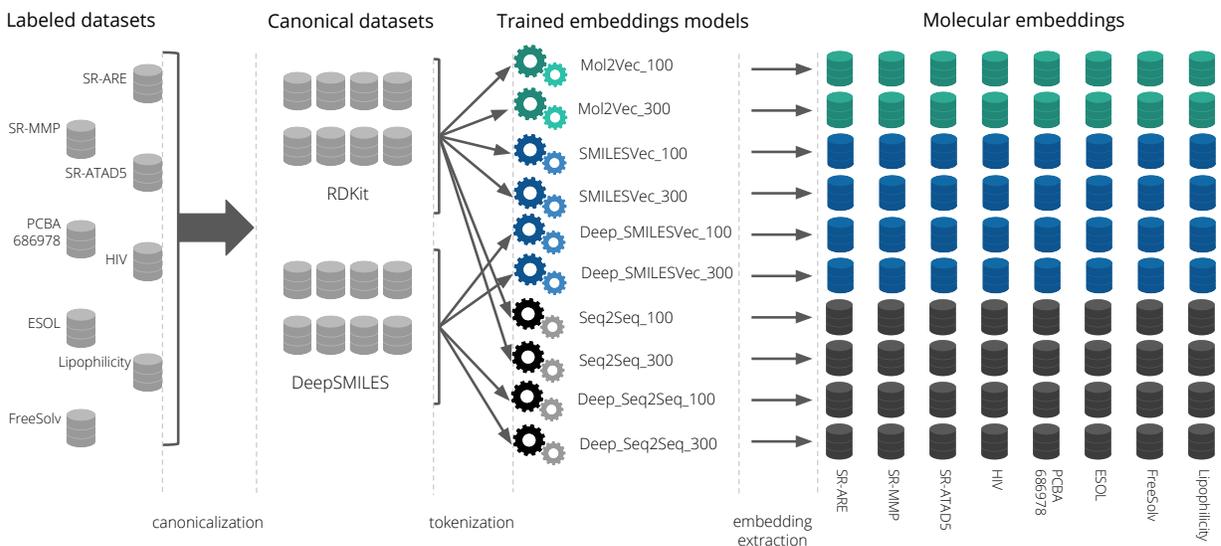
(b) Embedding extraction

Figure 3: Training and embedding extraction from unsupervised embedding methods: (a) The ZINC dataset was preprocessed and canonicalized following two different canonicalization procedures. The compounds in the two canonical datasets were tokenized as required by each of the unsupervised techniques, and then each unsupervised embedding model was trained. (b) The eight labeled datasets were canonicalized and tokenized according to the requirements of each method. Afterward, molecular embeddings were obtained for each labeled dataset from all the unsupervised embedding models trained in the previous step.
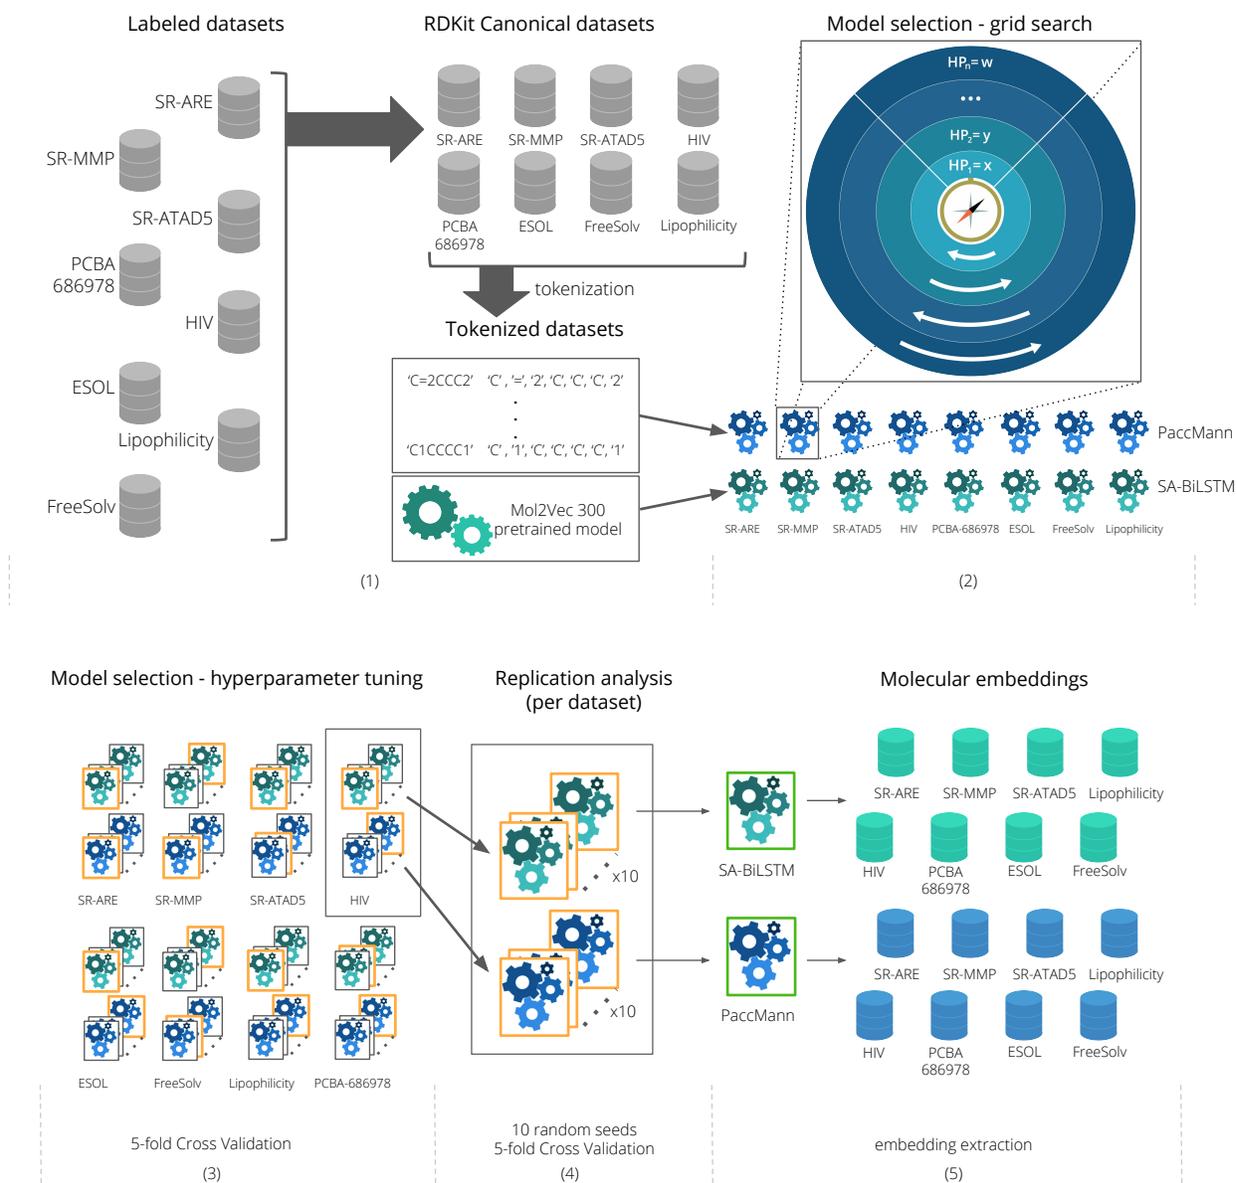
Figure 4: Training and embedding extraction from supervised methods: (1) The compounds in the eight labeled datasets were preprocessed and tokenized. (2) The first step of model selection on each labeled dataset consisted of an exploratory grid search for the best performing hyperparameter combination for each method. As a result, a subset of hyperparameter combinations was chosen for each labeled dataset. (3) In the second step of the model selection stage, each hyperparameter combination in the subset was tested on a five-fold cross-validation process. As a result, only one combination of hyperparameters was selected for each labeled dataset (highlighted in yellow). (4) Ten trials of the selected combination using different random seeds on five-fold cross-validation were trained. (5) An embedding model was trained using a stratified data partition and the embeddings for the eight labeled datasets were extracted from such model.

four different classifiers, whereas in the case of the regression tasks we trained three different regression models. In order to account for a fair comparison among all representations, we repeated the following steps for each representation:

1. For the classification tasks, we built and trained a Naïve Bayes classifier (NB), a Support Vector Machine (SVM) with an RBF kernel,[57] a Random Forest classifier (RF) and a shallow feed-forward neural network (FFNN). We used the Complement Naïve Bayes, SVM and RF implementations provided by Scikit-learn,[58] whereas the FFNNs were built and trained using Keras[59] and Tensorflow.[60] For the smallest datasets (*SR-ARE, SR-MMP* and *SR-ATAD5*), we trained single SVMs, whereas for the two largest labeled datasets (*HIV* and *PCBA-686978*), we employed the Scikit-learn implementation of a bagging classifier of ten SVMs, each trained on a stratified sample consisting of one-tenth of the dataset. For each classification method, we conducted a hyperparameter grid search using a stratified data partition. Full details about the hyperparameters tested on each classifier are provided in Appendix B.

2. In the case of the regression tasks, we trained a Ridge regression (Ridge), a Gradient Boosting Regressor (GBR) and a shallow feed-forward neural network (FFNN). We used the Ridge and GRB implementations provided by Scikit-learn,[58] and the FFNNs were built and trained using Keras[59] and Tensorflow.[60] We conducted a hyperparameter grid search on a stratified data partition for each regression model. All details about the hyperparameters tested on each model are provided in Appendix B.

3. We selected the best parameterizations of each method and ran a five-fold cross-validation training phase.

4. We selected the single best parameterization based on the results obtained during the five-fold cross-validation stage. Afterward, we ran ten trials of five-fold cross-validation using the same folds as before and using a different random seed for each run. This step

was carried out only for the stochastic methods (i.e., RF, GBR, FFNN and ensembles of SVMs).

During the training process we used the same folds for each labeled dataset in every stage of the experiments. In all cases, we measured the average results of the five validation folds and their 95% confidence intervals. The main reason behind this evaluation strategy is to account for a fair comparison of the results, minimizing the potential biases introduced by the variance of a particular random data partition, which can severely influence the results especially in small datasets, as it has been reported in previous studies.[16,61]

Results and discussion

This section discusses the results obtained for the five classification and three regression tasks defined by each of the labeled datasets described in Table 1. We compare the results using all the molecular embedding techniques on the four classification and three regression methods described in the previous section. We also report the results obtained by the supervised methods during the embedding learning stage, i.e., the *fitting* results, and the results obtained from traditional molecular representations (i.e. *molecular descriptors, ECFPs and MACCs keys*).

Throughout all stages of our experimental workflow, we measured the performance of our models using different metrics. The results on the regression tasks were evaluated by means of four metrics: *Root Mean Squared Error (RMSE), Mean Absolute Error (MAE) and Coefficient of Determination (R^2)*. In turn, the classification results were evaluated by eight metrics: *Sensitivity (Sn), Specificity (Sp), Precision, Accuracy (Acc), Balanced Accuracy (BA), F_1 score, H_1 score and Area Under the ROC Curve (AUC)*. *BA* is the arithmetic mean of *Sn* and *Sp* and F_1 score is the harmonic mean of *Sn* and *Precision*. To avoid overriding the notation for F_1 score, we hereby redefine the harmonic mean of *Sn* and *Sp* as H_1 score, which is also called F_1 or F score elsewhere in the literature.[62] We prioritized

$RMSE$ for hyperparameter selections in regression settings, and F_1 score and H_1 score in classification tasks. F_1 and H_1 scores are suitable in contexts of highly imbalanced datasets (such as in the case of HIV and SR-ATAD5) in contrast to metrics like Acc.[62,63] All the results obtained for each step of our experimental workflow in all eight labeled datasets can be found in the Supplementary Material.

In order to assess the statistical significance of the results, we conducted a series of statistical tests through which we compared the F_1 score and $RMSE$ results obtained by the different molecular representations and QSAR models on each dataset. We first conducted a two-way ANOVA where the molecular representation and the choice of the prediction model were considered the two independent variables. Upon eventually finding that the results of different molecular representations were significantly different, we conducted a *post-hoc* pairwise Tukey test[64] with a global confidence level of 95%. Full tables with the results of these statistical tests can be found in the Supplementary Material.

We accounted for the potential negative effects of class imbalance in classification tasks by three means. First, by training and evaluating the QSAR models using stratified partitions, in which the global proportions of active and inactive compounds (as described in Table 1) is present in each partition or fold. Second, by using weighed cost functions while training the supervised embedding models (*SA-BiLSTM* and *PaccMann*) and FFNNs, in a way such that the loss of the network during training is adjusted with a class weight proportional to the class imbalance. Third, by carefully selecting the metrics for model selection and evaluation of the results, as discussed above.

Our first research question (Q1) aimed at determining whether learned molecular representations could outperform traditional molecular representations in QSAR modeling. To answer this question, we compared the results obtained using the traditional molecular representations—*molecular descriptors*, *ECFPs* and *MACCS keys*—to the results obtained using both supervised and unsupervised molecular embeddings. After conducting the two-way ANOVA test, we proceeded to compare the fifteen representations under study, one

QSAR model at a time, using a pairwise Tukey test.

For the classification tasks, and as it can be seen in Figure 5, traditional representations were among the top-performing representations, followed by *SA-BiLSTM* and *Mol2Vec_300*. Traditional molecular representations yielded the best results for all datasets in NB, SVM and RF classifiers, which were significantly better than most learned embeddings. In the case of the imbalanced datasets—*SR-ATAD5* and *HIV*—, *ECFP* was significantly better than other representations using NB, and among the top-performing representation for the other classifiers, as shown in Figures 5 (c) and (d). *MACCS keys* and *molecular descriptors* also yielded significantly better results than learned embeddings in most datasets and classifiers. In the case of *PCBA-686978* (Figure 5 (e)), the best results were obtained by traditional representations in all classifiers except for FFNN. There were no significant differences among the results obtained using the three traditional representations. The best results for FFNN were generally obtained using *SA-BiLSTM*: this was observed for datasets *SR-ARE*, *SR-MMP* and *SR-ATAD5*. For dataset *HIV*, the best FFNN results were obtained by *ECFP*, showing a significant difference to all learned embeddings, whereas *Mol2Vec_300* yielded the best results on FFNN for dataset *PCBA-686978*. In all cases, these results were significantly better than those obtained using other learned embeddings, as shown in Figure 7. In general, the results using unsupervised embeddings did not match the results using traditional representations, except for *Mol2Vec* embeddings in FFNN classifiers. The statistical tests showed in all cases that the differences among these results were significant.

For regression tasks, as shown in Figure 6, a similar scenario is observed: traditional molecular representations were among the top performing representations in all datasets and for all regression methods. In particular, *molecular descriptors* and *SA-BiLSTM* generally attained the best results on GBR and FFNN in all three datasets. In particular for GBR, the differences between the results of these two representations were not significant. *SA-BiLSTM*, which was the top performing representation in FFNN, yielded significantly better results than those obtained by the unsupervised representations. A similar phenomenon was

observed for traditional representations, except for *Mol2Vec_100* and *Mol2vec_300* which generally yielded results on a par with the traditional representations. In dataset *Lipophilicity*, although the results for FFNN were different among the representations, these differences were no statistically significant among traditional representations, supervised embeddings, and unsupervised embeddings *Mol2Vec* and *SMILESVec*, as it can be seen in Figure 7 (c).

Our second research question (Q2) intended to determine whether supervised molecular embeddings could surpass unsupervised molecular embeddings in classification and regression tasks. We conducted a two-way ANOVA test among all learned molecular embeddings. Afterward, we performed a pairwise Tukey test to compare the results using the supervised embeddings against the unsupervised embeddings on each classifier.

As shown in Figures 5 and 6, the results obtained using the supervised representation *SA-BiLSTM* were generally significantly better than the results obtained employing unsupervised embeddings. In classification tasks, *SA-BiLSTM* was the top performing learned embedding in all balanced datasets (*SR-ARE*, *SR-MMP* and *PCBA-686978*). This observation, however, did not hold in the case of the NB-based classifiers, for which no significant differences among classification results were observed in any of the datasets. In the case of datasets *SR-ATAD5*, *HIV* and *PCBA-686978*, *SA-BiLSTM* attained similar results to those obtained by *Mol2Vec_100* and *Mol2Vec_300* in all classifiers except RF. This might be explained by the fact that *SA-BiLSTM* is trained using a form of *Mol2Vec* embeddings. In the case of regression tasks, *SA-BiLSTM* was also the top performing embedding in all datasets, showing significantly better results than those obtained by unsupervised representations. In dataset *Lipophilicity*, no significant differences were observed between the results of *PaccMann* and *SMILESVec* embeddings in GBR, and between the results of *Mol2Vec* and *SA-BiLSTM* embeddings in FFNN.

In the case of *PaccMann*, the results on the classification datasets did not show significant differences to the results obtained employing any of the unsupervised embeddings in the majority of the cases, except for FFNN-based classifiers, where the results are significantly
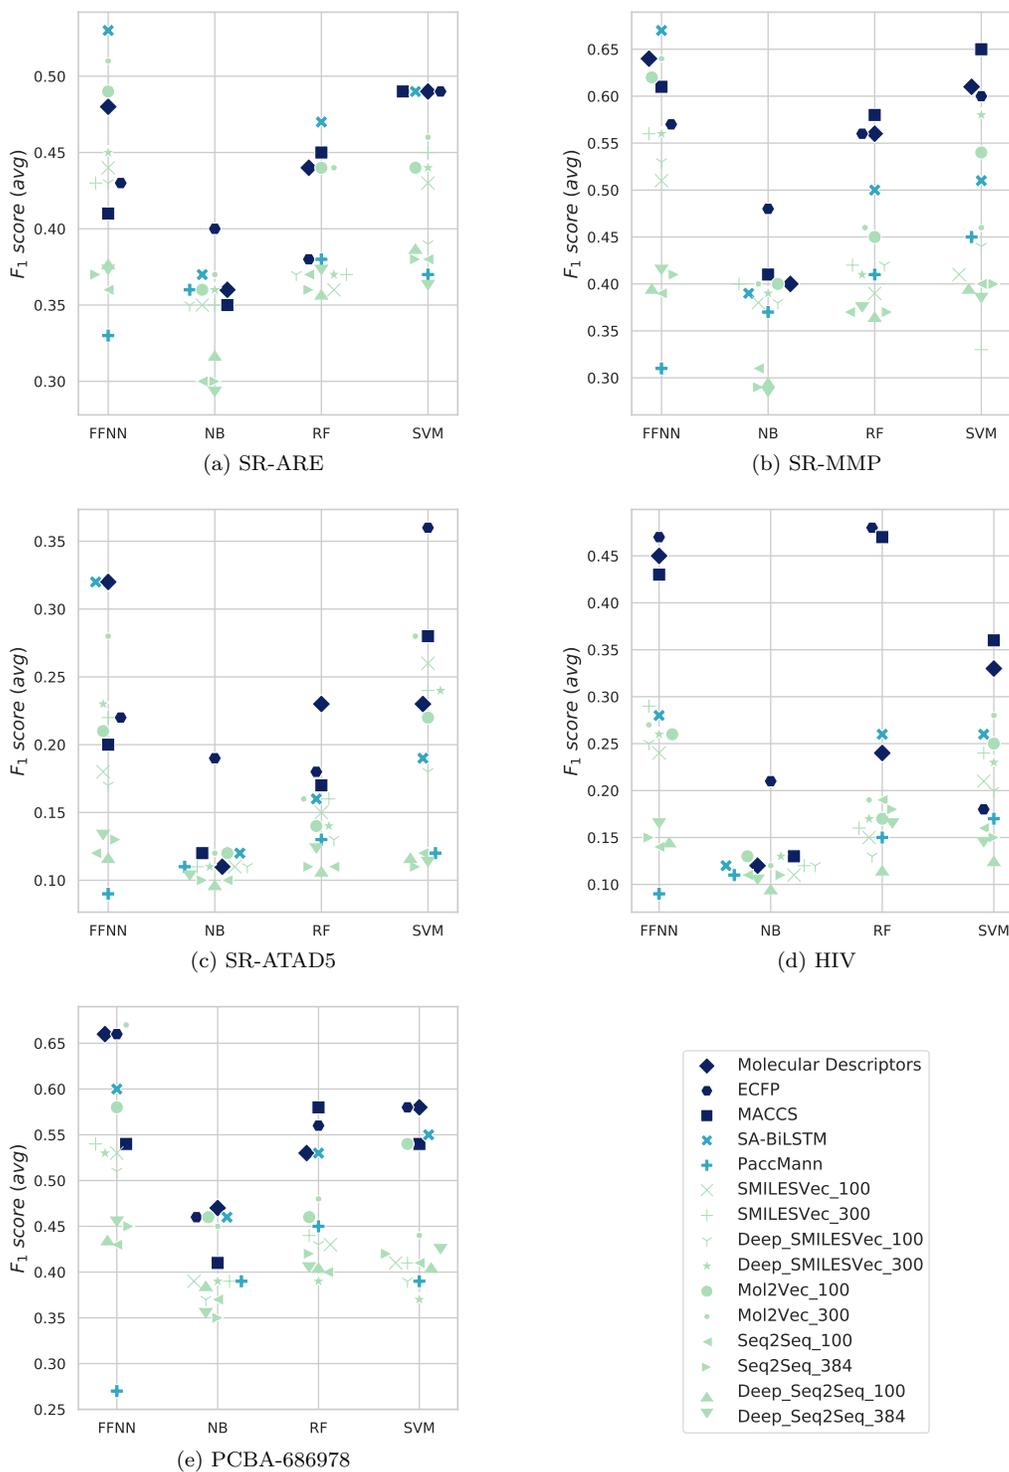
Figure 5: F_1 scores for the five classification datasets. Dark blue denotes traditional representations, light blue shows supervised embeddings, and green denotes unsupervised embeddings. Random horizontal jitter is applied to the markers to avoid overlap. Traditional representations attained the best performances in most cases, matched or followed by *SA-BiLSTM* and *Mol2Vec* embeddings.

28
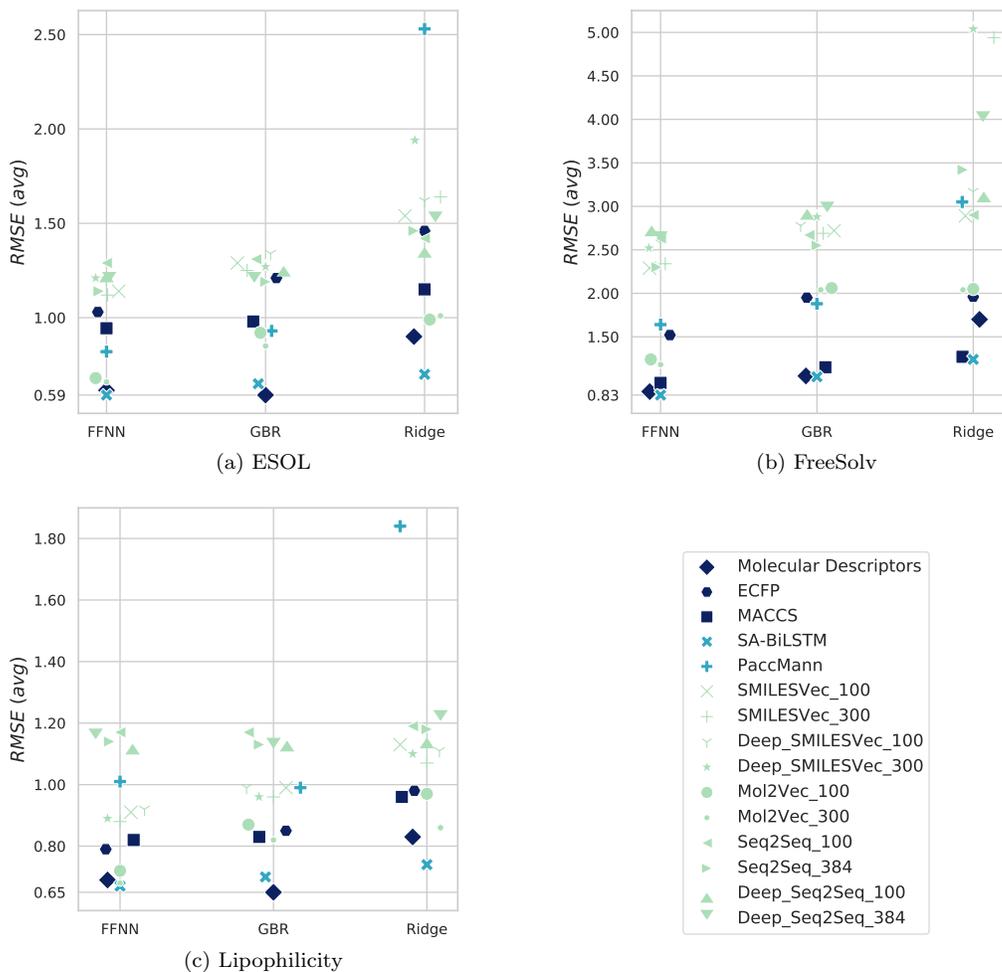
Figure 6: *RMSE* values for the three regression datasets. Dark blue denotes traditional representations, light blue shows supervised embeddings, and green denotes unsupervised embeddings. Random horizontal jitter is applied to the markers to avoid overlap. Traditional representations generally yielded the top results, matched or followed by *SA-BiLSTM* and *Mol2Vec* embeddings. In particular, *molecular descriptors* and *SA-BiLSTM* yielded the best results on FFNN and GBR in all three datasets.
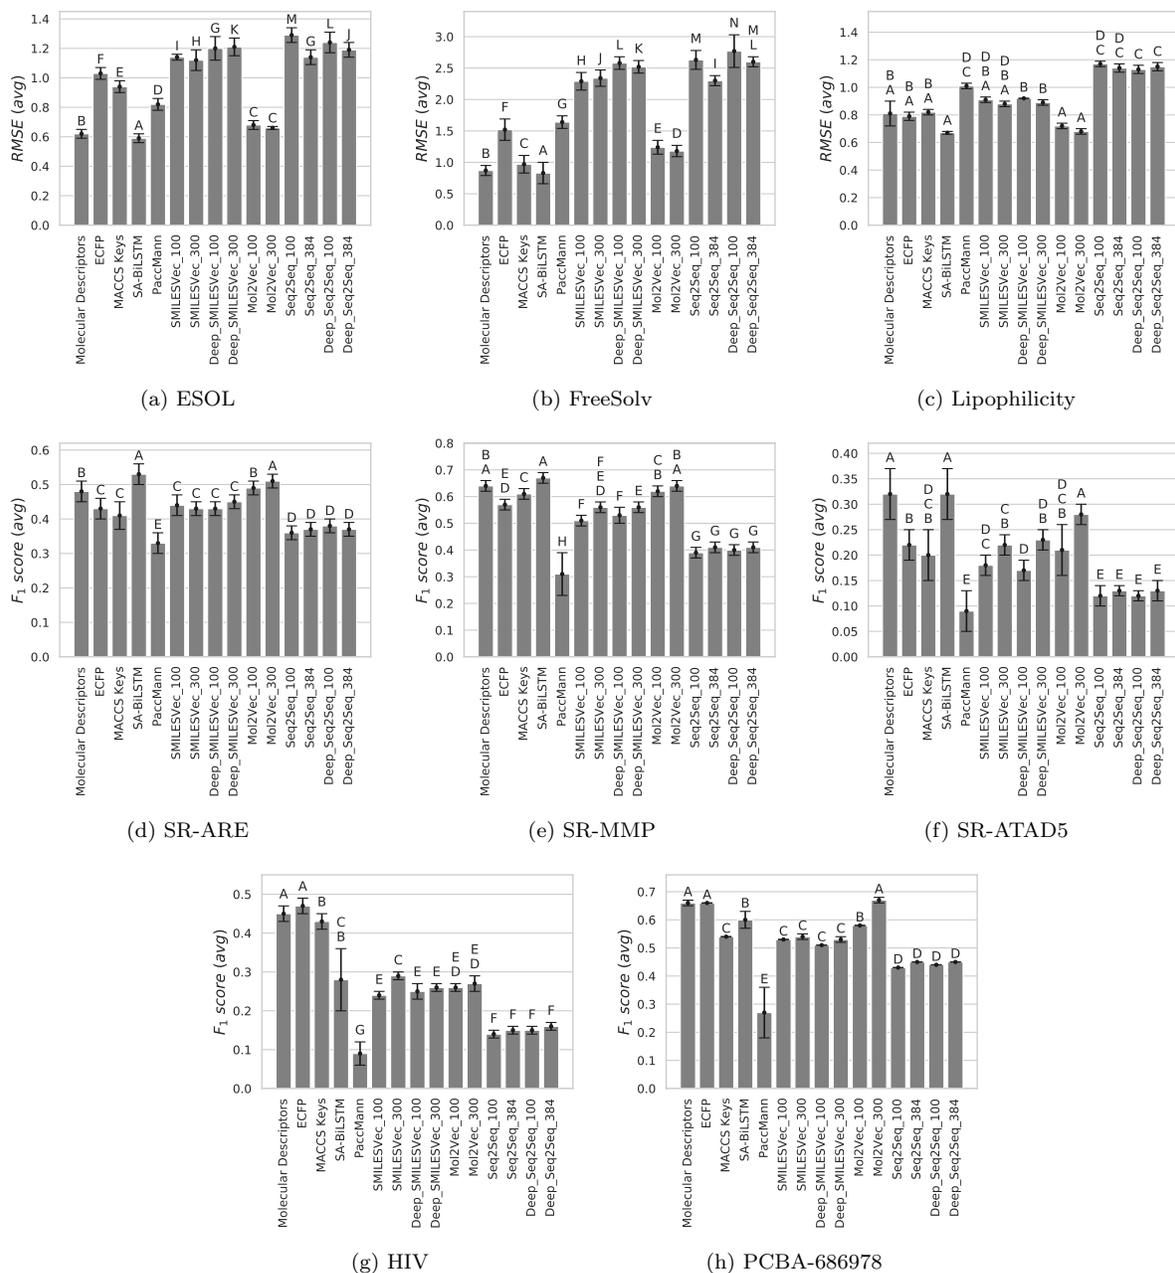
Figure 7: *RMSE* and F_1 scores obtained by all molecular representations in FFNN for regression (a, b, c) and classification datasets (d, e, f, g, h), respectively. Means with the same letter are not significantly different, according to the pairwise Tukey test. In the classification datasets, the results obtained using *Seq2Seq* embeddings are not statistically different in any of the datasets. *Mol2Vec* tends to show different results depending on their embedding size and are not significantly different to *SA-BiLSTM* embeddings or traditional representations. In datasets *ESOL* and *FreeSolv*, the results yielded by most representations were significantly different, whereas in dataset *Lipophilicity* there were no statistically significant differences among traditional representations, supervised embeddings, and unsupervised embeddings *Mol2Vec* and *SMILESVec*.

worse than all other representations. This can be observed in Figures 5 and 7. A different scenario is observed in the regression datasets, where *PaccMann* embeddings exhibited a significantly better performance than all unsupervised embeddings, except for *SMILESVec_100* and *Deep_SMILESVec_100* in *Lipophilicity*, and *Mol2Vec* in *ESOL* and *FreeSolv*, whose results in GBR were not significantly different. In Ridge regression, *PaccMann* results were significantly worse than those of the remaining representations.

When analyzing the unsupervised embeddings, for datasets *ESOL* and *FreeSolv* there were no significant differences between the results yielded by *SMILESVec* and *Seq2Seq* embeddings. In the case of the supervised embedding methods, as shown in Figure 8 and 9, *SA-BiLSTM* was significantly better than *PaccMann* for all datasets. It is worth noticing that the *fitting* results obtained by *PaccMann* in all datasets were significantly better than those obtained by the *PaccMann* embeddings using classifiers and regression methods. In the case of *SA-BiLSTM*, its *fitting* results were surpassed by at least one other classification technique in all datasets except for *PCBA-686978* and *Lipophilicity*.

Our third research question (Q3) aimed to determine whether the canonical form of the SMILES formulas used during training or the embedding sizes had a significant impact on the classification performance of QSAR models. To answer this question, we analyzed the three unsupervised embedding techniques separately. These results can be seen in Figures 5 and 6. We conducted separate two-way ANOVA tests and the corresponding pairwise Tukey tests to compare the results obtained by the following groups of molecular embeddings:

- *SMILESVec_100, SMILESVec_300, Deep_SMILESVec_100* and *Deep_SMILESVec_300*: No significant differences were observed among the results in any of the five classification datasets, except for SVM in the case of the balanced datasets (*SR-ARE*, *SR-MMP* and *PCBA-686978*), where all results were significantly different. We observed that the results of *SMILESVec_300* and *Deep_SMILESVec_300* tended to be significantly better than those obtained by *SMILESVEC_100* and *Deep_SMILESVec_100* when using FFNN for building the QSAR models on datasets *SR-MMP* and *SR-ATAD5*, as

31
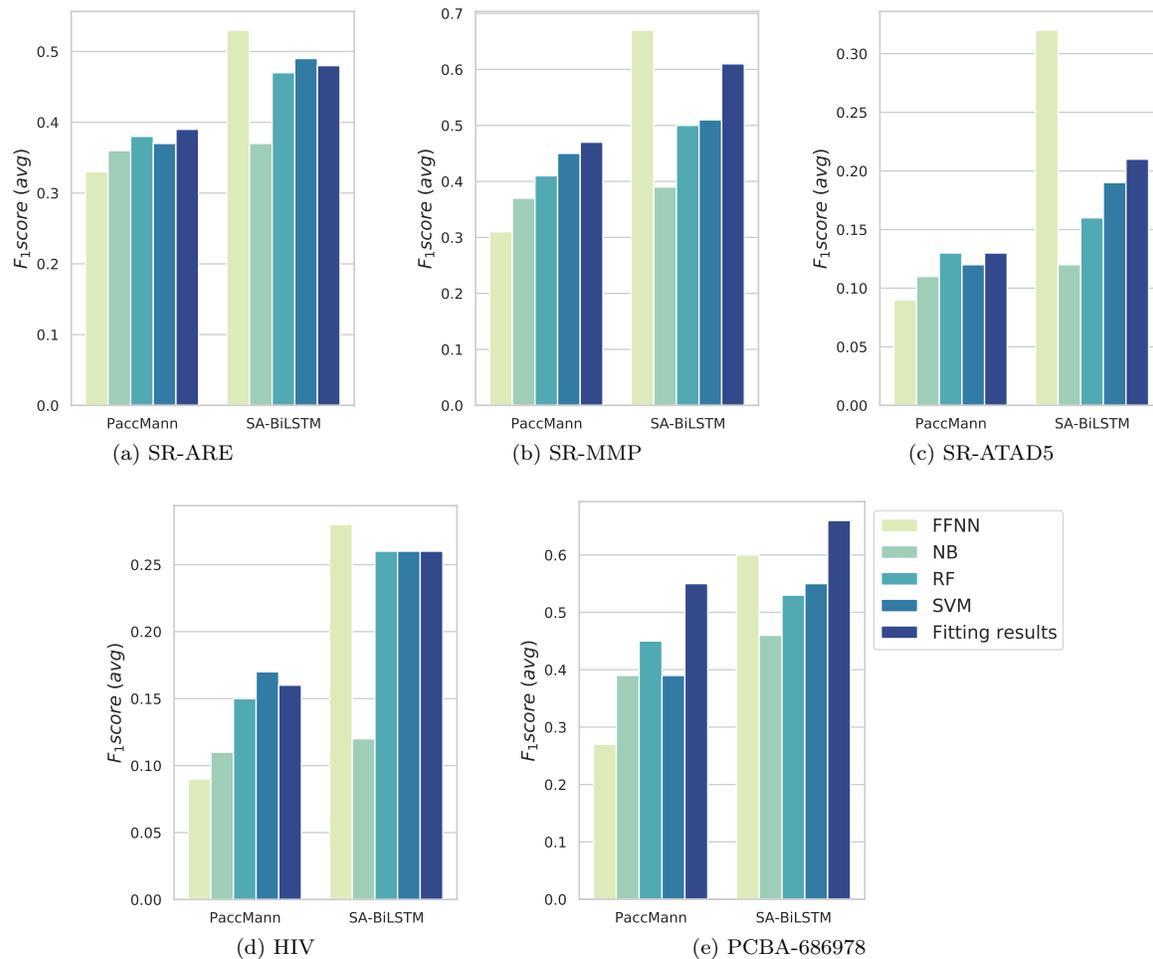
Figure 8: F_1 scores for the five classification datasets using supervised embeddings. *SA-BiLSTM* obtained top results among the supervised embeddings, especially on FFNN. While *PaccMann* results were significantly worse than those obtained by *SA-BiLSTM*, its *fitting* results—i.e., the classification results obtained from training the embedding model— were usually better than those obtained on other classifiers.
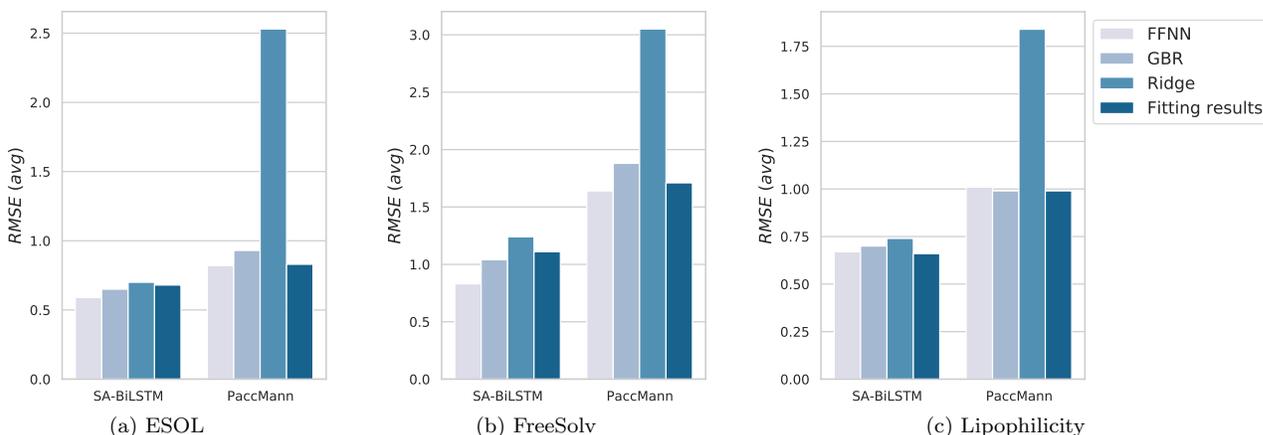
Figure 9: *RMSE* results for the three regression datasets using supervised embeddings. *SA-BiLSTM* obtained significantly better results than *PaccMann*. The *fitting* results of *PaccMann* surpassed all other regression results obtained with the same embedding.

it can be seen in Figures 5 and 7. In the case of regression tasks the differences were generally significant. Only for dataset *Lipophilicity* we found no significant differences when varying the canonical form of the SMILES formulas in GBR.

- *Mol2Vec_100* and *Mol2Vec_300*: The results obtained were significantly different when changing the size of the embedding from 100 to 300 throughout all datasets except for *FreeSolv*. As it happened in the case of *SMILESVec*, 300-dimension embeddings outperformed 100-dimension embeddings. This phenomenon was observed in all classifiers and regression models except for NB and Ridge, whose results were not significantly different.

- *Seq2Seq_100, Seq2Seq_384, Deep_Seq2Seq_100* and *Deep_Seq2Seq_384*: In the case of *Seq2Seq*, while the changes in the canonicalization or the size of embeddings did not yield significantly different results in any of the five classification datasets, significant differences were observed in the results for regression datasets *ESOL* and *FreeSolv* when varying the canonical form of SMILES formulas.

In addition to our comparative analysis, we present the results obtained by the top-performing molecular representations on each dataset side by side to the results reported in

the reference papers[14,17] and other state-of-the-art results for the datasets under study[15,16,47] in Tables 4 and 5. In spite of these tables, which we present in an effort to connect with other results reported in the literature, it is worth mentioning that a direct comparison among our results and the reference ones should not be made, considering that the vast majority of such papers report validation results on different sizes of fixed partitions.

The referenced papers reported regression performance in terms of $RMSE$ for datasets *ESOL, FreeSolv, Lipophilicity*, and classification performance in terms of AUC for datasets *SR-ARE, SR-MMP, SR-ATAD5, HIV*. In the case of dataset *PCBA-686978*, the results were reported in terms of Acc,[47] even though such metric is not appropriate for imbalanced datasets.[62,63] As shown in Table 4, the results of our experiments on the classification datasets are either on a par or surpass the results obtained by the reference papers on the internal validation sets, which arguably shows that our results are sound and that the extensive model selection stage was crucial. In the case of the regression tasks, shown in Table 5, we were able to match or surpass the results by one of the reference papers,[15] but not those reported in the other reference paper.[16] However, we reproduced their experiments on SVM and $XGBoost$ using the source code provided by the authors, and the results appear to be on a par with our results.

Finally, we present the classification results obtained by ten runs per molecular representation using FFNN classifiers, each using a different random initialization, in Figure 10. These results are expressed in terms of F_1 score and AUC, and they show the dispersion of the results of each molecular representation. Interestingly, while $ECFP$ attained the top results in terms of F_1 score for most classification datasets, $SA\text{-}BiLSTM$ matches or surpasses $ECFP$ when analyzing the AUC results in all of these datasets except for *PCBA-686978*. As it can be seen in Table 4 and in Figure 10, $SA\text{-}BiLSTM$ exhibits top AUC results on FFNN, often on a par with *molecular descriptors* and *Mol2Vec* embeddings. There is low dispersion within different runs using the same molecular representations, with the exception of $SR\text{-}ATAD5$, a fairly small and highly imbalanced dataset. When looking at the results

Table 4: Top results reported in the reference papers for the five classification datasets, in terms of *AUC* and *Acc*, and top results obtained through our experimental workflow marked with an asterisk (*). The best results per dataset are highlighted in **bold**. *na* denotes that the authors of the reference paper do not provide information about the confidence interval.

| Representation | Classifier | AUC | | | | Acc |
		SR-ARE	SR-MMP	SR-ATAD5	HIV	PCBA-686978
SA-BiLSTM[14]	fitting	$0.81 \pm na$	$0.90 \pm na$	$0.86 \pm na$	$0.81 \pm na$	-
Mol2Vec[17]	RF	0.83 ± 0.05	0.83 ± 0.05	0.83 ± 0.05	-	-
Weave[15]	fitting	$\mathbf{0.83 \pm 0.01}$	0.83 ± 0.01	0.83 ± 0.01	0.74 ± 0.04	-
ECFP[15]	XGBoost	0.78 ± 0.02	0.78 ± 0.02	0.78 ± 0.02	$\mathbf{0.84 \pm 0.00}$	-
SMILES-BERT[47]	fine-tuning	-	-	-	-	$\mathbf{0.88 \pm na}$
SA-BiLSTM (*)	FFNN	$\mathbf{0.83 \pm 0.02}$	$\mathbf{0.91 \pm 0.01}$	$\mathbf{0.87 \pm 0.02}$	0.83 ± 0.01	0.80 ± 0.04
Mol2Vec_300 (*)	FFNN	0.81 ± 0.01	0.88 ± 0.01	0.84 ± 0.01	0.81 ± 0.01	0.83 ± 0.00
ECFP (*)	FFNN	0.74 ± 0.02	0.84 ± 0.01	0.70 ± 0.04	0.78 ± 0.01	0.87 ± 0.00

Table 5: Top results reported in the reference papers for the three regression datasets, in terms of *RMSE*, and top results obtained through our experimental workflow marked with an asterisk (*). The best results per dataset are highlighted in **bold**. *na* denotes that the authors of the reference paper do not provide information about the confidence interval.

| Representation | Regression Model | RMSE | | |
		ESOL	FreeSolv	Lipophilicity
Attentive FP[16]	fitting	$\mathbf{0.48 \pm na}$	$0.52 \pm na$	$\mathbf{0.52 \pm na}$
MOE + FPs[16]	SVM	$0.62 \pm na$	$\mathbf{0.42 \pm na}$	$0.55 \pm na$
MOE + FPs[16]	XGBoost	$0.51 \pm na$	$0.69 \pm na$	$\mathbf{0.52 \pm na}$
MPNN[15]	fitting	0.55 ± 0.02	1.20 ± 0.02	0.76 ± 0.03
Weave[15]	fitting	0.57 ± 0.04	1.19 ± 0.08	0.73 ± 0.01
GC[15]	fitting	1.05 ± 0.15	1.35 ± 0.15	0.68 ± 0.04
SA-BiLSTM (*)	FFNN	0.59 ± 0.03	0.83 ± 0.17	0.67 ± 0.01
Mol2Vec_300 (*)	FFNN	0.66 ± 0.01	1.18 ± 0.09	0.68 ± 0.02
Molecular descriptors (*)	FFNN	0.62 ± 0.03	0.87 ± 0.08	0.69 ± 0.03
Molecular descriptors (*)	GBR	0.59 ± 0.04	1.05 ± 0.12	0.65 ± 0.02

in terms of F_1 score, traditional representations surpassed the majority of unsupervised embeddings in the case of the largest datasets, *HIV* and *PCBA-686978*. For *PCBA-686978*, *Mol2Vec_300* obtained the top results using FFNN. When analyzing other metrics, such as *Precision* and *Sp* (available in the Supplementary Material), *ECFP* attained better results than *SA-BiLSTM* and *Mol2Vec_300*, which explains its superiority in terms of F_1 score.

Based on our observations we can conclude that, in general, learned molecular embed-

(a) SR-ARE

(b) SR-MMP

(c) SR-ATAD5

(d) HIV

(e) PCBA-686978
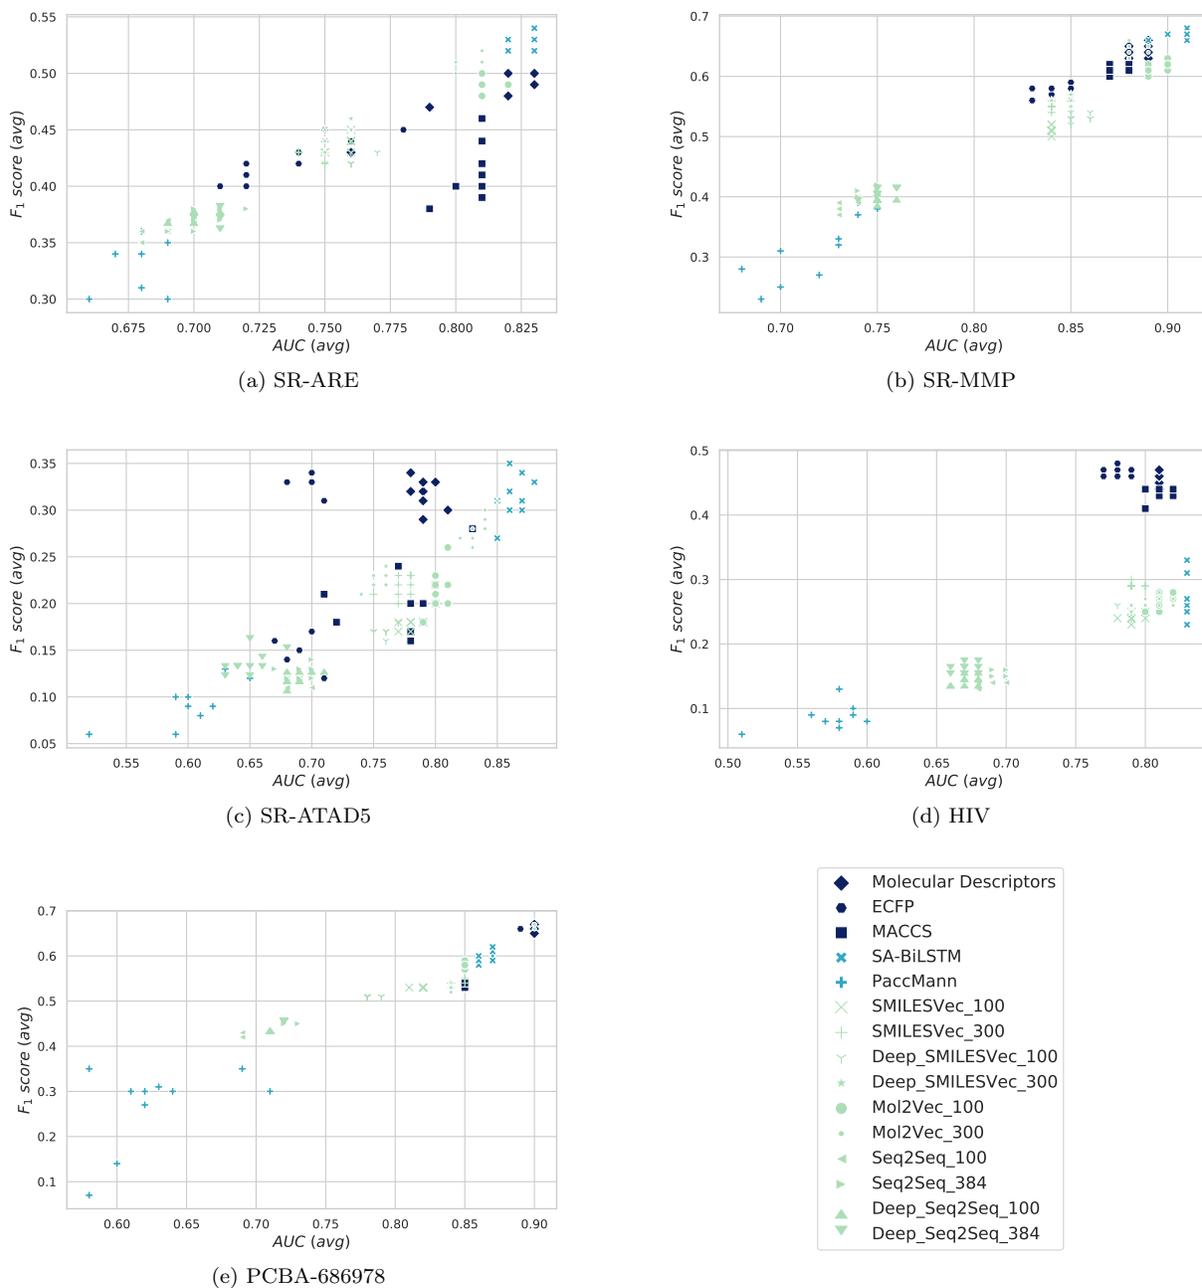
Figure 10: Results in terms of F_1 score and AUC for ten runs of FFNN classifiers, each using a different random initialization. *SA-BiLSTM* obtained the best results in most classification datasets, often matched by *molecular descriptors* or *Mol2Vec* embeddings. Traditional representations yielded better F_1 score results than *SA-BiLSTM* in the largest datasets (*HIV* and *PCBA-686978*).

dings did not broadly surpass the results obtained by traditional molecular representations. Also, most of the unsupervised embeddings methods did not match the results obtained by traditional molecular representations. Among the unsupervised embedding techniques, *Mol2Vec* yielded the best results, often performing significantly better than the results obtained using *SMILESVec* or *Seq2Seq*. This could be related to the preprocessing step on the SMILES formulas, which is based on the algorithm for computing *ECFP* fingerprints, in contrast to the simple tokenization steps applied to SMILES formulas in the case of the other two techniques. *Seq2Seq* obtained the lowest performance in terms of F_1 score and $RMSE$ of all representations, especially in the case of imbalanced classification datasets *SR-ATAD5* and *HIV*.

In the case of the supervised embeddings, on the one hand, *SA-BiLSTM* yielded the best results among the learned representations, which were often on a par with those obtained using traditional molecular representations. On the other hand, *PaccMann* did not yield good results in classification tasks. This could be explained by the high dimensionality of *PaccMann* embeddings—ranging from 6976 to 24, 200 dimensions, as shown in Table 3—, which is often a cause of traditional machine learning models having difficulties in learning patterns from the data and, consequently, of a low predictive performance.[65,66] It is also worth noticing that the original *PaccMann* embedding model proposed by Oskooei et al.[13] was developed as a multimodal encoder for anticancer compound sensitivity prediction. Another possible reason for the significant difference between the results of the two supervised embedding techniques might be that *SA-BiLSTM* consists of a multi-head self-attention model, whereas *PaccMann* is a single-head self-attention model.

Our experiments were designed to attain the best possible performance for each molecular representation. We experimented with different classification and regression techniques, dataset sizes and class imbalance scenarios. We also tested different parameterizations and variations of the input data for each embedding method. The conclusions derived in this paper are presented after observing consistent and sound results throughout our analysis,

which were also similar to other results reported in the literature.[14–17,47] This pattern of learned molecular representations only matching the performance of traditional representations, while unexpected, can also been observed in other studies that experimented with learned molecular embeddings for QSAR modeling.[16–22] However, such results are not conclusive, as they either are not supported by any significance test, or there was not a systematic comparison where hyperparameters were tuned fairly.We argue that our results prove the importance of conducting a thorough and careful experimental comparison of molecular embedding techniques and the potential role of learned representations in QSAR modeling.

Despite our findings, learned molecular embeddings might be suitable for many other tasks than QSAR modeling in the spectrum of drug design and virtual screening, which opens a wide range of experimental possibilities and future work.[5,6,27] While traditional representations are computed following a standard algorithm and looking at a single molecule at a time, learned embeddings can be computed from large sets of compounds, potentially yielding richer representations that could be suitable for molecular similarity analysis.[5,67] Besides, techniques like self-attention might yield good embeddings for tasks like molecular substructure search, molecular docking,[5,6,14] or for linking molecular substructures to bioactivity profiles.[14] Another possible direction would be to combine traditional representations with molecular embeddings, thus leveraging their strengths.

Conclusions

Many novel deep learning-based algorithms for learning molecular representations have been proposed in recent years. However, along with the proliferation of novel algorithms, it is crucial to conduct methodical and careful comparisons between different existing algorithms in order to shed light onto the key aspects of the embedding process. In this paper, we carried out an extensive comparison, which consisted of over $25,000$ trained models, using five different techniques for learning molecular representations. We tested their suitability

for QSAR modeling in five classification and three regression tasks, considering aspects such as embedding size, SMILES canonicalization and deep learning algorithms, and compared them against traditional molecular representations.

As a result of our experimental analysis, we found no evidence of molecular embeddings surpassing molecular descriptors or fingerprints, which might indicate that traditional representations are as good as molecular embeddings are at QSAR prediction tasks. We also found that unsupervised molecular embeddings exhibited lower performance than supervised techniques in general. Despite these results that contradict the current trend in molecular representations, drug discovery has numerous flourishing areas of research where there is room for learned representations, such as *de novo* drug design, molecular docking or virtual screening.[5,6] Contrary to deterring the community from using and developing new algorithms for learning molecular representations, we hope that the results observed in this paper can serve as an incentive to design new ways of using molecular embeddings to leverage their potential for different tasks in drug discovery.

Appendices

Appendix A: hyperparameters tested during grid search of supervised embedding methods

- *PaccMann*: in the self-attention layer, we varied the number of hidden units (i.e. 16, 50 and 100) and the attention depth value (i.e. 20, 50, 100 and 256). In the feed-forward neural network used for property prediction, we varied the number of hidden units (i.e. $[512, 256, 64, 16]$, $[100, 50, 20, 5]$, $[150, 50, 10]$ and $[100, 20, 5]$) and the activation functions (i.e. *ReLU* and *sigmoid*)[68] of the dense layers. We trained with and without a sinusoidal positional encoding function applied to the inputs,[12] and with and without a weighed loss function based on the imbalance in each class.[69] We varied the regularization hyperparameters such as dropout coefficients per layer (i.e. 0.5, 0.25

and 0.15) and *lambda* values (i.e. 0.0001, 0.005 and 0.001) for L2-regularization.[70] The learning rate was fixed to 0.001, and we used Adam optimizer.[71] We used a minibatch size of 512 for the small dataset (i.e. *SR-ARE, SR-MMP, SR-ATAD5*) and a minibatch size of 2048 for the large datasets (i.e. *HIV, PCBA-686978*). As a result of this grid search process, we tested a total of 175 different combinations of hyperparameters for each labeled dataset.

- *SA-BiLSTM*: since each run of this method takes a significant amount of time to train, we decided to focus on the hyperparameters tested by the authors in the original paper. We varied the number of attention heads (i.e. 5, 10, 15 and 18) and the attention depth value (i.e. 10, 20, 50 and 100) in the self-attention layer, as well as the number of hidden units in the Bi-LSTM RNN (i.e. 64 and 128). Following the values indicated in the reference paper, we used a minibatch size of 64 for all datasets and used a gradient clipping coefficient of 0.3. The dropout coefficient per layer was set to 0.2 and the λ value for L2-regularization[70] to 0.01. The models were trained for a maximum of 1000 epochs, as per the reference paper. We trained with and without applying a weighed loss function which compensated the class imbalance. As a result of this grid search stage, we tested a total of 64 different combinations of hyperparameters for each labeled dataset.

Appendix B: hyperparameters tested during grid search of classifiers and regression models

For each classification method, we conducted a hyperparameter grid search using a stratified data partition. In the case of the SVMs, we varied the regularization coefficient $c \in \{0.01, 0.05, 0.1, 0.25, 0.5, 1, 5, 10, 20\}$, and for the RF classifier we varied the maximum depth of the tree $md \in \{2, 3, 5, 8, 10, 20\}$. For the FFNNs, we varied the number of nodes in each layer (i.e. $[150, 50, 10]$, $[100, 50, 10]$ and $[100, 20, 5]$), the L2-regularization hyperparame-

ter $\lambda \in \{0.0001, 0.005, 0.001\}$, the minibatch size $b \in \{64, 128, 256, 512\}$, the activation function used in the hidden layers (i.e. ReLU and $tanh$) and the early-stopping $patience$ coefficient $p \in \{70, 100, 200, 500\}$. We also tested different learning rates $\alpha \in \{0.00001, 0.0001, 0.001\}$.

In the case of the regression models, for Ridge regression we varied the regularization coefficient $\alpha \in \{0.001, 0.05, 0.1, 0.5, 1, 10\}$, the tolerance $tol \in \{0.00001, 0.0001, 0.001\}$, and set the maximum number of iterations for convergence $max_iter = 10000$. For Gradient Boosting Regressor (GBR), we varied the number of estimators $n_e stimators \in \{50, 100, 200, 500\}$, the maximum depth of the tree $max_depth \in \{2, 3, 5, 810, 20\}$ and the minimum number of samples required to split an internal node $min_samples_split \in \{2, 3\}$. Finally, for the FFNNs we varied the number of nodes in each layer (i.e. $[200, 70, 15]$, $[150, 50, 10]$, $[100, 50, 10]$ and $[100, 20, 5]$), the L2-regularization hyperparameter $\lambda \in \{0.0001, 0.005, 0.001, 0.01, 0.1\}$, the activation function used in the hidden layers (i.e. ReLU and $tanh$), the early-stopping $patience$ coefficient $p \in \{100, 200, 500\}$ and the learning rate value $\alpha \in \{0.00001, 0.0001, 0.001\}$. We set the minibatch size $b = 200$, and also tested two loss functions: $RMSE$ and MSE.

Data availability statement

In order to ensure the reproducibility of our experimental workflow, all resources, materials and source code used in this paper are either properly cited or we have made them publicly available. The data, trained models and source code underlying this article are available in a public repository [xiii].

Competing interests

The authors declare no competing interests.

[xiii] https://csunseduar-my.sharepoint.com/:f:/g/personal/virginia_sabando_cs_uns_edu_ar/EjUkG4X2A31EgJ0Aj0EjveYBMcooO8mKIpQoHquoQtdUhw

Funding

Acknowledgments

References

(1) Wu, Z.; Zhu, M.; Kang, Y.; Leung, E. L.-H.; Lei, T.; Shen, C.; Jiang, D.; Wang, Z.; Cao, D.; Hou, T. Do we need different machine learning algorithms for QSAR modeling? A comprehensive assessment of 16 machine learning algorithms on 14 QSAR data sets. *Briefings in Bioinformatics* **2020**,

(2) Wu, Z.; Jiang, D.; Hsieh, C.-Y.; Chen, G.; Liao, B.; Cao, D.; Hou, T. Hyperbolic relational graph convolution networks plus: a simple but highly efficient QSAR-modeling method. *Briefings in Bioinformatics* **2021**,

(3) Cherkasov, A.; Muratov, E. N.; Fourches, D.; Varnek, A.; Baskin, I. I.; Cronin, M.;

[xiv]https://www.deepsense.ca/.
[xv]https://www.ace-net.ca/.
[xvi]https://www.calculquebec.ca/.
[xvii]https://www.computecanada.ca/.

Dearden, J.; Gramatica, P.; Martin, Y. C.; Todeschini, R., et al. QSAR modeling: where have you been? Where are you going to? *Journal of Medicinal Chemistry* **2014**, *57*, 4977–5010.

(4) Todeschini, R.; Consonni, V. *Molecular descriptors for chemoinformatics: volume I: alphabetical listing/volume II: appendices, references*; John Wiley & Sons, 2009; Vol. 41.

(5) Chuang, K. V.; Gunsalus, L. M.; Keiser, M. J. Learning Molecular Representations for Medicinal Chemistry: Miniperspective. *Journal of Medicinal Chemistry* **2020**, *63*, 8705–8722.

(6) Elton, D. C.; Boukouvalas, Z.; Fuge, M. D.; Chung, P. W. Deep learning for molecular design—a review of the state of the art. *Molecular Systems Design & Engineering* **2019**, *4*, 828–849.

(7) Chen, H.; Engkvist, O.; Wang, Y.; Olivecrona, M.; Blaschke, T. The rise of deep learning in drug discovery. *Drug Discovery Today* **2018**, *23*, 1241–1250.

(8) Bouhedjar, K.; Boukelia, A.; Khorief Nacereddine, A.; Boucheham, A.; Belaidi, A.; Djerourou, A. A natural language processing approach based on embedding deep learning from heterogeneous compounds for quantitative structure–activity relationship modeling. *Chemical Biology & Drug Design* **2020**, *96*, 961–972.

(9) David, L.; Thakkar, A.; Mercado, R.; Engkvist, O. Molecular representations in AI-driven drug discovery: a review and practical guide. *Journal of Cheminformatics* **2020**, *12*, 1–22.

(10) Weininger, D. SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules. *Journal of Chemical Information and Computer Sciences* **1988**, *28*, 31–36.

(11) Wu, Z.; Pan, S.; Chen, F.; Long, G.; Zhang, C.; Philip, S. Y. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems* **2020**, *32*, 950–957.

(12) Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; Polosukhin, I. Attention Is All You Need. *arXiv e-prints* **2017**, arXiv:1706.03762.

(13) Oskooei, A.; Born, J.; Manica, M.; Subramanian, V.; Sáez-Rodríguez, J.; Rodríguez Martínez, M. PaccMann: Prediction of anticancer compound sensitivity with multimodal attention-based neural networks. *arXiv e-prints* **2018**, arXiv:1811.06802.

(14) Zheng, S.; Yan, X.; Yang, Y.; Xu, J. Identifying structure–property relationships through SMILES syntax analysis with self-attention mechanism. *Journal of Chemical Information and Modeling* **2019**, *59*, 914–923.

(15) Wu, Z.; Ramsundar, B.; Feinberg, E. N.; Gomes, J.; Geniesse, C.; Pappu, A. S.; Leswing, K.; Pande, V. MoleculeNet: a benchmark for molecular machine learning. *Chemical Science* **2018**, *9*, 513–530.

(16) Jiang, D.; Wu, Z.; Hsieh, C.-Y.; Chen, G.; Liao, B.; Wang, Z.; Shen, C.; Cao, D.; Wu, J.; Hou, T. Could graph neural networks learn better molecular representation for drug discovery? A comparison study of descriptor-based and graph-based models. *Journal of cheminformatics* **2021**, *13*, 1–23.

(17) Jaeger, S.; Fulle, S.; Turk, S. Mol2vec: unsupervised machine learning approach with chemical intuition. *Journal of Chemical Information and Modeling* **2018**, *58*, 27–35.

(18) Gómez-Bombarelli, R.; Wei, J. N.; Duvenaud, D.; Hernández-Lobato, J. M.; Sánchez-Lengeling, B.; Sheberla, D.; Aguilera-Iparraguirre, J.; Hirzel, T. D.; Adams, R. P.; Aspuru-Guzik, A. Automatic chemical design using a data-driven continuous representation of molecules. *ACS Central Science* **2018**, *4*, 268–276.

(19) Goh, G. B.; Siegel, C.; Vishnu, A.; Hodas, N. O.; Baker, N. Chemception: A Deep Neural Network with Minimal Chemistry Knowledge Matches the Performance of Expert-developed QSAR/QSPR Models. *arXiv e-prints* **2017**, arXiv:1706.06689.

(20) Yang, K.; Swanson, K.; Jin, W.; Coley, C.; Eiden, P.; Gao, H.; Guzman-Perez, A.; Hopper, T.; Kelley, B.; Mathea, M., et al. Analyzing learned molecular representations for property prediction. *Journal of Chemical Information and Modeling* **2019**, *59*, 3370–3388.

(21) Yang, K.; Swanson, K.; Jin, W.; Coley, C.; Eiden, P.; Gao, H.; Guzman-Perez, A.; Hopper, T.; Kelley, B.; Mathea, M.; Palmer, A.; Settels, V.; Jaakkola, T.; Jensen, K.; Barzilay, R. Correction to Analyzing Learned Molecular Representations for Property Prediction. *Journal of Chemical Information and Modeling* **2019**, *59*, 5304–5305, PMID: 31814400.

(22) Chithrananda, S.; Grand, G.; Ramsundar, B. ChemBERTa: Large-Scale Self-Supervised Pretraining for Molecular Property Prediction. *arXiv e-prints* **2020**, arXiv:2010.09885.

(23) Sterling, T.; Irwin, J. J. ZINC 15–ligand discovery for everyone. *Journal of Chemical Information and Modeling* **2015**, *55*, 2324–2337.

(24) Sabando, M. V.; Ulbrich, P.; Selzer, M.; Byška, J.; Mičan, J.; Ponzoni, I.; Soto, A. J.; Ganuza, M. L.; Kozlíková, B. ChemVA: Interactive Visual Analysis of Chemical Compound Similarity in Virtual Screening. *IEEE Transactions on Visualization and Computer Graphics* **2021**, *27*, 891–901.

(25) Cereto-Massagué, A.; Ojeda, M. J.; Valls, C.; Mulero, M.; Garcia-Vallvé, S.; Pujadas, G. Molecular fingerprint similarity search in virtual screening. *Methods* **2015**, *71*, 58–63.

(26) Grisoni, F.; Consonni, V.; Todeschini, R. *Computational Chemogenomics*; Springer New York: New York, NY, 2018; pp 171–209.

(27) Schneider, G. Virtual screening: an endless staircase? *Nature Reviews Drug Discovery* **2010**, *9*, 273–276.

(28) Rogers, D.; Hahn, M. Extended-connectivity fingerprints. *Journal of Chemical Information and Modeling* **2010**, *50*, 742–754.

(29) Durant, J. L.; Leland, B. A.; Henry, D. R.; Nourse, J. G. Reoptimization of MDL keys for use in drug discovery. *Journal of Chemical Information and Computer Sciences* **2002**, *42*, 1273–1280.

(30) Seth, A.; Roy, K. QSAR modeling of algal low level toxicity values of different phenol and aniline derivatives using 2D descriptors. *Aquatic Toxicology* **2020**, *228*, 105627.

(31) Yang, L.; Wang, Y.; Chang, J.; Pan, Y.; Wei, R.; Li, J.; Wang, H. QSAR modeling the toxicity of pesticides against Americamysis bahia. *Chemosphere* **2020**, *258*, 127217.

(32) Gao, K.; Nguyen, D. D.; Sresht, V.; Mathiowetz, A. M.; Tu, M.; Wei, G.-W. Are 2D fingerprints still valuable for drug discovery? *Physical Chemistry Chemical Physics* **2020**, *22*, 8373–8390.

(33) Sabando, M. V.; Ponzoni, I.; Soto, A. J. Neural-based approaches to overcome feature selection and applicability domain in drug-related property prediction. *Applied Soft Computing* **2019**, *85*, 105777.

(34) Liu, S.; Furkan Demirel, M.; Liang, Y. N-Gram Graph: Simple Unsupervised Representation for Graphs, with Applications to Molecules. *arXiv e-prints* **2018**, arXiv:1806.09206.

(35) Swann, E.; Sun, B.; Cleland, D.; Barnard, A. Representing molecular and materials data for unsupervised machine learning. *Molecular Simulation* **2018**, *44*, 905–920.

(36) Öztürk, H.; Ozkirimli, E.; Özgür, A. A novel methodology on distributed representations of proteins using their interacting ligands. *Bioinformatics* **2018**, *34*, i295–i303.

(37) Xu, Z.; Wang, S.; Zhu, F.; Huang, J. Seq2seq fingerprint: An unsupervised deep molecular embedding for drug discovery. Proceedings of the 8th ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics. 2017; pp 285–294.

(38) Kuzminykh, D.; Polykovskiy, D.; Kadurin, A.; Zhebrak, A.; Baskov, I.; Nikolenko, S.; Shayakhmetov, R.; Zhavoronkov, A. 3d molecular representations based on the wave transform for convolutional neural networks. *Molecular Pharmaceutics* **2018**, *15*, 4378–4385.

(39) Shi, T.; Yang, Y.; Huang, S.; Chen, L.; Kuang, Z.; Heng, Y.; Mei, H. Molecular image-based convolutional neural network for the prediction of ADMET properties. *Chemometrics and Intelligent Laboratory Systems* **2019**, *194*, 103853.

(40) Özçelik, R.; Öztürk, H.; Özgür, A.; Ozkirimli, E. ChemBoost: A chemical language based approach for protein-ligand binding affinity prediction. *arXiv e-prints* **2018**, arXiv:1811.00761.

(41) Mikolov, T.; Chen, K.; Corrado, G.; Dean, J. Efficient Estimation of Word Representations in Vector Space. *arXiv e-prints* **2013**, arXiv:1301.3781.

(42) Segler, M. H.; Kogej, T.; Tyrchan, C.; Waller, M. P. Generating focused molecule libraries for drug discovery with recurrent neural networks. *ACS Central Science* **2018**, *4*, 120–131.

(43) Popova, M.; Isayev, O.; Tropsha, A. Deep reinforcement learning for de novo drug design. *Science Advances* **2018**, *4*, eaap7885.

(44) Cho, K.; van Merrienboer, B.; Gulcehre, C.; Bahdanau, D.; Bougares, F.; Schwenk, H.; Bengio, Y. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. *arXiv e-prints* **2014**, arXiv:1406.1078.

(45) Joshi, C. Transformers are Graph Neural Networks. `https://thegradient.pub/transformers-are-gaph-neural-networks/`, 2020.

(46) Hochreiter, S.; Schmidhuber, J. Long short-term memory. *Neural computation* **1997**, *9*, 1735–1780.

(47) Wang, S.; Guo, Y.; Wang, Y.; Sun, H.; Huang, J. SMILES-BERT: large scale unsupervised pre-training for molecular property prediction. Proceedings of the 10th ACM International Conference on Bioinformatics, Computational Biology and Health Informatics. 2019; pp 429–436.

(48) Devlin, J.; Chang, M.-W.; Lee, K.; Toutanova, K. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv e-prints* **2018**, arXiv:1810.04805.

(49) Lipinski, C. A. Lead- and drug-like compounds: the rule-of-five revolution. *Drug Discovery Today: Technologies* **2004**, *1*, 337–341.

(50) Landrum, G. RDKit: open-source cheminformatics http://www. rdkit. org. 2016.

(51) Delaney, J. S. ESOL: estimating aqueous solubility directly from molecular structure. *Journal of chemical information and computer sciences* **2004**, *44*, 1000–1005.

(52) Mobley, D. L.; Guthrie, J. P. FreeSolv: a database of experimental and calculated hydration free energies, with input files. *Journal of computer-aided molecular design* **2014**, *28*, 711–720.

(53) Bento, A. P.; Gaulton, A.; Hersey, A.; Bellis, L. J.; Chambers, J.; Davies, M.; Krüger, F. A.; Light, Y.; Mak, L.; McGlinchey, S., et al. The ChEMBL bioactivity database: an update. *Nucleic acids research* **2014**, *42*, D1083–D1090.

(54) Dalke, A. DeepSMILES: An Adaptation of SMILES for Use in Machine-Learning of Chemical Structures. 2018.

(55) Schwaller, P.; Gaudin, T.; Lanyi, D.; Bekas, C.; Laino, T. "Found in Translation": predicting outcomes of complex organic chemistry reactions using neural sequence-to-sequence models. *Chemical Science* **2018**, *9*, 6091–6098.

(56) Moriwaki, H.; Tian, Y.-S.; Kawashita, N.; Takagi, T. Mordred: a molecular descriptor calculator. *Journal of Cheminformatics* **2018**, *10*, 1–14.

(57) Schölkopf, B.; Tsuda, K.; Vert, J.-P. *Kernel methods in computational biology*; MIT press, 2004.

(58) Pedregosa, F. et al. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* **2011**, *12*, 2825–2830.

(59) Chollet, F. Keras. 2015; https://github.com/fchollet/keras, online July 2020.

(60) Abadi, M. et al. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. 2015; `http://tensorflow.org/`, Software available from tensorflow.org.

(61) Baumann, D.; Baumann, K. Reliable estimation of prediction errors for QSAR models under model uncertainty using double cross-validation. *Journal of cheminformatics* **2014**, *6*, 1–19.

(62) Sokolova, M.; Japkowicz, N.; Szpakowicz, S. Beyond accuracy, F-score and ROC: a family of discriminant measures for performance evaluation. Australasian Joint Conference on Artificial Intelligence. 2006; pp 1015–1021.

(63) Chawla, N. V.; Bowyer, K. W.; Hall, L. O.; Kegelmeyer, W. P. SMOTE: synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research* **2002**, *16*, 321–357.

(64) Tukey, J. W., et al. *Exploratory data analysis*; Reading, Mass., 1977; Vol. 2.

(65) Goodarzi, M.; Dejaegher, B.; Vander Heyden, Y. Feature selection methods in QSAR studies. *Journal of AOAC International* **2012**, *95*, 636–51.

(66) Alsenan, S.; Al-Turaiki, I.; Hafez, A. Autoencoder-based Dimensionality Reduction for QSAR Modeling. 2020 3rd International Conference on Computer Applications & Information Security (ICCAIS). 2020; pp 1–4.

(67) Huang, B.; Von Lilienfeld, O. A. Communication: Understanding molecular representations in machine learning: The role of uniqueness and target similarity. 2016.

(68) Sharma, S. Activation functions in neural networks. 2017.

(69) Domingos, P. Metacost: A general method for making classifiers cost-sensitive. Proceedings of the Fifth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. 1999; pp 155–164.

(70) Moody, J.; Hanson, S.; Krogh, A.; Hertz, J. A. A simple weight decay can improve generalization. *Advances in Neural Information Processing Systems* **1995**, *4*, 950–957.

(71) Kingma, D. P.; Ba, J. Adam: A Method for Stochastic Optimization. *arXiv e-prints* **2014**, arXiv:1412.6980.