

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/344012873>

ChemVA: Interactive Visual Analysis of Chemical Compound Similarity in Virtual Screening

Preprint · August 2020

CITATIONS

0

READS

34

9 authors, including:



Matias Selzer

Universidad Nacional del Sur

9 PUBLICATIONS 14 CITATIONS

[SEE PROFILE](#)



Jan Mičan

Masaryk University

2 PUBLICATIONS 6 CITATIONS

[SEE PROFILE](#)



Ignacio Ponzoni

Universidad Nacional del Sur

133 PUBLICATIONS 585 CITATIONS

[SEE PROFILE](#)



M. Luján Ganuza

Universidad Nacional del Sur

17 PUBLICATIONS 36 CITATIONS

[SEE PROFILE](#)

Some of the authors of this publication are also working on these related projects:



Augmented Reality in Geology [View project](#)



Visualization of geological data [View project](#)

ChemVA: Interactive Visual Analysis of Chemical Compound Similarity in Virtual Screening

María Virginia Sabando*, Pavol Ulbrich*, Matías Selzer, Jan Byška, Jan Mičan, Ignacio Ponzoni, Axel J. Soto, María Luján Ganuza, Barbora Kozlíková




Fig. 1: Overview of the ChemVA interface: a) Hexagonal view with an overview of a selected 2D projection, b) Detail view showing all data items, which are colored according to a selected feature, c) 3D view enabling the user to observe the structural similarities of selected compounds, d) Table view listing other important features of the compounds. All views are interactively linked.

Abstract— In the modern drug discovery process, medicinal chemists deal with the complexity of analysis of large ensembles of candidate molecules. Computational tools, such as dimensionality reduction (DR) and classification, are commonly used to efficiently process the multidimensional space of features. These underlying calculations often hinder interpretability of results and prevent experts from assessing the impact of individual molecular features on the resulting representations. To provide a solution for scrutinizing such complex data, we introduce ChemVA, an interactive application for the visual exploration of large molecular ensembles and their features. Our tool consists of multiple coordinated views: Hexagonal view, Detail view, 3D view, Table view, and a newly proposed Difference view designed for the comparison of DR projections. These views display DR projections combined with biological activity, selected molecular features, and confidence scores for each of these projections. This conjunction of views allows the user to drill down through the dataset and to efficiently select candidate compounds. Our approach was evaluated on two case studies of finding structurally similar ligands with similar binding affinity to a target protein, as well as on an external qualitative evaluation. The results suggest that our system allows effective visual inspection and comparison of different high-dimensional molecular representations. Furthermore, ChemVA assists in the identification of candidate compounds while providing information on the certainty behind different molecular representations.

Index Terms—Virtual screening, visual analysis, dimensionality reduction, coordinated views, cheminformatics.

1 INTRODUCTION

*These authors contributed equally.

- *M. V. Sabando, I. Ponzoni, and A. J. Soto are with the Institute for Computer Science and Engineering (UNS-CONICET) and with the Department of Computer Science and Engineering, Universidad Nacional del Sur, Bahía Blanca, Argentina. E-mails: virginia.sabando@cs.uns.edu.ar, ip@cs.uns.edu.ar, axel.soto@cs.uns.edu.ar.*
- *P. Ulbrich, J. Byška, and B. Kozlíková are with the Visitlab, Faculty of Informatics, Masaryk University, Brno, Czech Republic. E-mails: paloulbrich@gmail.com, jan.byska@gmail.com, kozlikova@fi.muni.cz.*
- *M. N. Selzer and M. L. Ganuza are with the Institute for Computer Science and Engineering (UNS-CONICET) and with the VyGLab Research Laboratory (UNS-CICPBA), Department of Computer Science and Engineering, Universidad Nacional del Sur, Bahía Blanca, Argentina. E-mails: matias.selzer@cs.uns.edu.ar, mlg@cs.uns.edu.ar.*
- *J. Mičan is with the Loschmidt Laboratories, Department of Experimental Biology and RECETOX, and with the Faculty of Medicine, Masaryk University, Brno, Czech Republic. E-mail: honzamicann@gmail.com.*

Small organic chemical compounds are the cornerstone of drug design. New medications are found by exploring a large number of candidate compounds or by designing new ones. In the last decades, high-throughput screening has been the main procedure applied during the early stages of the drug discovery process [19, 41]. This process requires chemical synthesis, experimental testing of large libraries of compounds against a biological target (protein), and it has a high attrition rate, which makes the process costly and time-consuming. These drawbacks stimulated the development of virtual screening methods, which consist of computational techniques for identifying compounds binding to a drug target. Virtual screening allows to significantly narrow down the number of drug candidate compounds at a faster pace

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxxx/TVCG.201x.xxxxxx

while lowering costs [39, 90]. These reasons make virtual screening an essential part of the early-stage drug discovery process.

Computational techniques involved in virtual screening enable to simulate and test the fitness of the candidate compound towards the desired function without the need for chemical synthesis and expensive wet laboratory work [71]. They typically consist of a high-dimensional vector-based abstraction of a given molecule, which involves dealing with the challenges inherent in high-dimensional data [27]. Many visualization tools for virtual screening have been proposed in the past years to help medicinal chemists deal with the complex computational methods. Such tools have been focused on providing proper means to explore the chemical space and to enhance the interpretability of results so that trusted and explainable decisions can be made [20, 55]. A common strategy followed by such tools is the application of dimensionality reduction (DR) techniques [16, 55], which allow mapping compounds from a high-dimensional chemical space into a lower-dimensional (often 2D or 3D) representation [16]. However, visual exploration of the chemical space in the context of virtual screening entails a series of new challenges, not yet addressed by the existing tools.

One of these important yet unsolved challenges is to study compound similarity under the premise that similar molecules tend to have similar bioactivity profiles [26]. Visualization tools should enable the domain expert to interact with different sources of molecular information, and provide complementary views that help find similarity determinants. Most existing tools rely on a single mapping based on an arbitrarily chosen set of molecular features. A visualization tool for virtual screening should also provide views and interactions that allow the domain expert to assess the relationship between the bioactivity profiles of the analyzed compounds and their spatial organization in the projected low-dimensional space. Lastly, as a consequence of DR, pairwise distances between compounds in the low-dimensional mapping can differ from the corresponding pairwise distances in the high-dimensional space. Information about the trustworthiness of these mappings should be visually presented to the domain expert in an interpretable way. Despite its importance, most of the existing tools do not display this information.

These issues were the main driving force behind our research and, as a result, we introduce ChemVA, an interactive system for the visual exploration of chemical compounds, targeted for virtual screening. ChemVA provides domain experts with several linked views. One set of views, consisting of 2D plots, supports the visual inspection of multiple molecular representations after undergoing DR. Each DR method produces a projection that can be displayed and mutually compared using our newly proposed *Difference view*. Our tool handles the overplotting problem of such projections, which is common for large datasets. ChemVA also incorporates a correlation encoding to the plots, which conveys the trustworthiness of a low-dimensional projection based on the distortion with regard to the pairwise distances in the original space. The plots are complemented by an interactive *Table view*, which allows for sorting and filtering and shows detailed information about the compounds being displayed, focusing primarily on features related to drug-likeness. The tool also contains a *3D view* that allows to explore the structural similarity among selected compounds. This view performs an alignment of selected compounds with respect to their maximum common substructure, which serves to identify commonalities and differences among them. Finally, ChemVA allows loading new compounds to an existing dataset, which enables to study potential new drug candidates in the context of the dataset under study. The main contributions of this paper can be summarized as follows:

- A novel visualization tool for virtual screening that allows to compare and contrast multiple dimensionality reduction projections conveying complementary sources of molecular information.
- An approach for visually assessing the trustworthiness of each DR projection, which enables the user to focus on the most suitable vector-based molecular representations.
- A newly proposed *Difference view* that enables the comparison of multiple DR projections and assessment of their trustworthiness in terms of neighborhood preservation.

- Support for *de novo* drug design by means of a set of coordinated views that allows to analyze newly designed compounds and compare them with an existing set of chemical compounds.
- Functional validation of the proposed tool by means of two case studies and a qualitative testing of the proposed views. These evaluation activities were conducted by three domain experts and one visualization expert.

2 RELATED WORK

In this section, we conduct a brief survey of existing approaches in the different areas related to our work. We review DR techniques and their application to the visualization of high-dimensional spaces, focusing mainly on parametric DR methods. Then, we discuss the state-of-the-art molecular visualization and visual exploration tools for large molecular ensembles.

2.1 Visual Exploration of Multidimensional Data

Over the last few decades, a variety of visualization methods for multidimensional data have been proposed [5, 40, 82]. Discovering patterns in multidimensional data using a combination of visual and machine learning techniques represents a well-known challenge in visual analytics. Visual exploration of multidimensional data allows for the injection of unique human perceptual and cognitive abilities directly into the process of discovering multidimensional patterns [32]. The challenge mostly lies in having a large number of variables and their relationships that have to be considered simultaneously. As the number of variables increases, the user's ability to understand interactions and correlations between them is severely limited [80].

In this context, a variety of approaches have been introduced to visually convey high-dimensional data by using two-dimensional projections, so that salient structures or patterns can be perceived while exploring the projected data [45, 61, 91]. Therefore, one typical approach is to transform the original dataset using a DR technique, and then to visually encode only the reduced data [50, 65]. A frequently used visual encoding technique for showing the projected low dimensional data is a scatter plot. Sedlmair et al. [66] carried out an extensive investigation on the effectiveness of visual encoding choices, including 2D scatter plots, interactive 3D scatter plots, and scatter plot matrices. Their findings suggest that the 2D scatter plot is the most suitable approach to explore the output of different DR algorithms.

Many DR techniques have been proposed throughout the last decades to transform multidimensional feature spaces onto a low-dimensional manifold—typically onto a two or three-dimensional space—while trying to preserve neighborhood relationships among the original data instances [78, 86]. Among these methods, *t-distributed Stochastic Neighbor Embedding* (t-SNE) [77] has been widely adopted, specifically for the visualization of very high-dimensional data [43, 57, 70, 85]. Many of the proposed DR techniques are non-parametric, i.e., they find patterns in specific manifolds and do not provide means to map new data points from the original high-dimensional space to the latent space. Therefore, their out-of-sample extension is not possible. Linear methods for DR, such as PCA [84], LDA [12], or MSR [69], have been extensively used, but are constrained to linear transformations of the original data. In contrast to other non-linear techniques, e.g., SOMs [31], MDS [9] or autoencoders [34], which allow to map non-linear data, t-SNE tackles the *crowding problem* effectively by using a heavy-tailed distribution for arranging data points in the low dimensional space [77]. Practically speaking, addressing the crowding problem yields better-looking visualizations. Although t-SNE is non-parametric, there is a parametric t-SNE version [76], which seeks to overcome this limitation. Nonetheless, it is often difficult to find an optimal configuration of hyperparameters for these models, which in turn yields noisy projections compared to those obtained using the non-parametric version of the algorithm [44, 87, 93].

In particular, visualization tools for virtual screening should allow medicinal chemists to analyze how a previously unseen compound interacts with or relates to other known compounds. For this reason, DR techniques used in the visualization tools for virtual screening

should provide means for mapping unseen compounds and they should be fast enough to enable interactive use.

2.2 Visualization of Molecular Structures

Molecular visualization is one of the oldest branches of visualization as it has a well-established basis of visual representations, which has been largely embraced by biochemists and biologists. These—mostly 3D—representations are integrated into many available molecular visualization tools, such as PyMOL [64], VMD [21], Chimera [54], or YASARA [35]. Besides these tools, there are also several others—mainly web-based ones—that can be embedded into other applications; for instance, Jmol [25], JSmol [17], or 3Dmol.js [58]. However, none of these tools is directly applicable to the problem of visual exploration of large ensembles of molecular structures and the similarities in their structure and properties. A more extensive overview of the currently available approaches to molecular visualization and molecular systems can be found in the survey by Kozlíková et al. [33].

Visual inspection of small molecules is supported in different ways by some available tools. Smaller molecules are commonly visualized in 2D in order to better depict their structure, such as in LigPlot+ [37]. This tool serves mainly for exploring the interactions between a ligand and its target protein, by projecting its molecular structure and the protein amino acids onto the 2D plane. However, such a tool is not applicable to the task of analyzing multiple molecular compounds.

The exploration of large sets of chemical compounds constitutes an ideal task from the visual analytics viewpoint, and several such tools have already been proposed. The closest to our proposed tool is CheS- Mapper [16]. It enables to load data from several chemical databases, to calculate molecular descriptors, to observe clusters of compounds in a reduced 2D space, and to display the 3D representation of the compounds. Nonetheless, the tool does not provide an option to interactively compare different projections and does not support inspecting the physical-chemical properties of a newly added compound.

A similar tool is Data Warrior [62], which consists of several basic visualization methods and specialized views to show an overview of the whole chemical and pharmacophore space of the input dataset. However, the tool is not designed for specific tasks related to the early-stage drug discovery process that is addressed in our tool. More recently, Yoshimori et al. [89] proposed a technique for visually summarizing information about multiple Structure-Activity Relationship Matrices, based on molecular grid maps and 3D activity landscapes. Nevertheless, due to its specificity, it cannot be extended to operate with additional physical-chemical properties and descriptors.

Janssen et al. [24] introduced a method that uses a visual mapping based on t-SNE projections for finding new potential kinase inhibitors. The views provided by the tool show the results of clustering and a tree-like structure of compounds. However, such views are not designed to easily extract information about the structural similarity of compounds. A similar concept was introduced by Probst and Reymond [56]. Their approach is tailored to process very large high-dimensional datasets and the compound similarity is expressed by the proximity of compounds through tree branches.

Naveja and Medina-Franco [49] introduced the constellation graphs of chemical compounds clustered according to a shared core scaffold in the 2D projections of their geometry. Each cluster is annotated with a small view containing the projected core. While such constellation graphs provide an innovative approach for the visualization of compounds, they do not fit our requirements for visual exploration in the context of virtual screening. Synergy Maps [38] is another web-based application for displaying relationships between compounds and understanding potential synergies between them. The tool shows pairwise combinations of properties of compounds using a network representation. Nonetheless, the node-link diagrams do not scale properly, which hinders its application for large ensembles of compounds.

Although methods for assessing the trustworthiness of DR techniques have been proposed [2, 10, 53], to the best of our knowledge none of the existing approaches for visualization of molecular structures allows to perform this task. Moreover, they do not provide means to compare the results of multiple projections.

3 BACKGROUND

In this section, we provide details on the theoretical background of our proposed tool, which includes the molecular representations used and their semantics, as well as a short description of the chemical features of major relevance for drug-likeness.

3.1 Vector-Based Molecular Representations

Research in cheminformatics suggests that there are several different factors involved in assessing chemical compound similarity, which go beyond having similar geometric arrangements of atoms and bonds [3, 42]. Given the relevance of assessing compound similarity in virtual screening, it is important to provide medicinal chemists with a variety of vector-based molecular representations that capture different aspects of the compounds under study and that are complementary to each other. In this sense, ChemVA provides four different sources of data:

- *Extended Connectivity Fingerprints* (ECFPs) [59] encode the topological information of the molecular structure as a fixed-length vector of bits.
- *Daylight Fingerprints* [11] are path-based vectors of bits that encode all fragments of a molecule, which are obtained by traversing its molecular graph.
- *Molecular Descriptors* [74] are numerical values associated with the chemical constitution of a molecule. They convey information such as molecular weight, electronic configuration, solubility, etc.
- *Molecular Embeddings* [23] are a rather novel type of vector-based representations. They are typically obtained by applying machine learning models that are trained to learn data representations of molecules from large and diverse sets of compounds.

We selected these four vector-based molecular representations as our sources of data because they capture different information about the compounds, thus allowing for a broader analysis of the data under study, and because each of them is considered to be the state-of-the-art of their type. While we are utilizing these four vector-based molecular representations, it is worth noting that ChemVA remains independent from the chosen sources of data. Details of their calculation are provided in the Supplementary Material.

3.2 Molecular Features Related to Drug-Likeness

Some molecular features are of major interest to drug designers and medicinal chemists in the context of virtual screening. In particular, molecular features related to drug-likeness allow scrutinizing the compounds in terms of their viability as potential new drugs. ChemVA takes into consideration a subset of such features:

- *Molecular weight* represents the average mass of the molecule expressed in Daltons (Da). Drug-like compounds are expected to weigh from 160 to 500 Da.
- *logP* is a quantitative measure of lipophilicity, and indicates how easily a substance is absorbed by living tissue. Drug-like compounds often exhibit logP values from -0.4 to $+5.6$.
- *Acidic Dissociation Constant* and *Basic Dissociation Constant* are quantitative measures of the strength of an acid (or base) in a chemical solution. They are good indicators for a drug's ability to enter the bloodstream and to accumulate in tissues or secretions.
- *Lipinski's Rule of 5 (RO5)* is a rule of thumb to determine whether a compound is likely to be active in humans. This rule states that an orally active drug violates no more than one of four criteria based on threshold values of specific chemical features.
- *Weighed Quantitative Estimate of Drug-likeness (QED)* is a score computed based on chemical features linked to drug-likeness, and it ranges from 0 (*non drug-like*) to 1 (*drug-like*).

4 REQUIREMENTS

Over the course of one year, we conducted numerous sessions with the group of protein engineers from the Loschmidt Laboratories at Masaryk University. Based on their input, we identified several limitations of the existing approaches for virtual screening and summarized them into a list of requirements. The experts agreed that these requirements cover the most critical aspects of a virtual screening workflow. One of the experts, who is also the co-author of this paper, has dedicated significant time ensuring that our implementation addresses the requirements by iteratively checking and commenting on the progress. His 4-year research experience in protein engineering, along with his internal consultations with the head of his group, provided the necessary domain knowledge for the design of the tool.

R1: Overview and detailed analysis of a molecular ensemble in the low-dimensional space. For large datasets, scatter plots, which are commonly used to represent the DR output, may suffer from occlusion problems for large datasets. Therefore, the tool should provide visual support for the analysis of data on different levels of abstraction, from the overall distribution of the compounds within the 2D space to the detailed view of individual compounds for a selected region of interest.

R2: Visual inspection of multiple projections. A set of compounds can be expressed by different vector-based molecular representations, each yielding a different DR projection. The tool should enable the user to intuitively combine information encoded in the individual projections to allow studying at once the similarity between compounds based on different molecular representations. More specifically, this includes exploring similarities and differences between chemical compounds, expressed by the different projections. Additionally, the visual representations and interactions should help the domain experts evaluate the suitability of the selected DR model.

R3: Evaluation of the trustworthiness of projections. Users need proper visual support for assessing the trustworthiness of a low-dimensional projection based on the distortion with regard to pairwise distances between compounds in the original space. When such trustworthiness can be compared on different DR projections, users can focus the exploration on a subset of molecular representations.

R4: Comparison of compounds according to features related to drug-likeness. Chemical compounds have many features and descriptors related to drug-likeness that can complement other molecular representations in the task of virtual screening. Therefore, it is desirable to provide users with an option to visualize these additional features along with the projections.

R5: Comprehensible viewing of 3D structural similarity. The tool should support the inspection of individual compounds in terms of their 3D geometry, as well as the visual comparison of common 3D substructures in a selected set of compounds. For multiple compounds, such a view should convey the information about similarities and differences in their 3D structure.

R6: Possibility to add new compounds and comparison with the existing data. The tool should support the process of exploration of different features and bioactivity for newly added compounds. The new compound should be projected using the DR model and integrated into the remaining views, so that the user can compare its features to those of the compounds in the existing dataset.

5 CHEMVA DESIGN AND IMPLEMENTATION

The design process of ChemVA followed the requirements listed in Section 4. These requirements are addressed through several coordinated views, which pose challenges to the layout of the tool. ChemVA was developed in JavaScript using D3.js v5 [6] and a Node.js server [73] built using the Express framework and REST API for the connection to supporting services. We used Unity3d [72] ported to WebGL [29] for the development of 3D components. Two back-end web services, used for computing the alignment of 3D structures and for computing the 2D coordinates of newly added compounds, were developed using Flask web framework [60] and written in Python v3. Excluding one of the views [13], all of the functionality was designed and written in-house,

prioritizing the system responsiveness and real-time experience with interactions.

In this section, we first describe the individual views supported by ChemVA in detail, then we briefly introduce the DR model applied in our tool, and lastly we present the main functionality supporting the analysis of newly added compounds.

5.1 ChemVA Views

The initial layout, depicted in Figure 1, contains the Hexagonal view, Detail view, and 3D view (Figure 1 A, B, and C, respectively) in the top row of the canvas, whereas the bottom row contains the Table view (Figure 1 D). These views account for requirements R1, R4 and R5. For tasks involving the comparison of projections, stated in requirement R2, additional views need to be displayed. After careful consideration, we decided to position these views between the two rows of the basic layout (see Figure 8), as they need to be as close as possible to the default 2D plot views. This row can be shown or collapsed on demand and allows the user to choose which views should be included (e.g., Hexagonal, Detail, or Difference views). The purpose of this layout is to allow the user to observe the data from different perspectives. In particular, the Difference view shows the differences between the compound neighborhoods from one projection into another, thus allowing the user to compare different projections. The tool uses the color scheme proposed by Okabe & Ito [22], so it is accessible to users with color vision deficiency.

5.1.1 2D Plot Views

The core of ChemVA consists of 2D plots, which give the user an overview of the distribution of compounds in a selected DR projection using a given molecular representation. This overview is supported by the *Hexagonal view*, a well-adopted and commonly used approach to visualize the outcome of DR techniques [66]. The Hexagonal view aims to overcome the overplotting problem, in which the projected data items overlap and cause visual clutter, thus limiting the interpretability, especially for datasets evincing high similarity between data items. The Hexagonal view of ChemVA seeks to solve this problem by aggregating individual data items into individual hexagons. The user can interactively select a subset of hexagons of interest and explore the distribution of individual data items within the *Detail view*. The combination of the *Hexagonal view* and the *Detail view* aims to fulfill requirement R1. Finally, since ChemVA is tailored to support the visual comparison of different projections, it also offers the *Difference view*. This view was specifically designed to address that task, which is stated in requirements R2 and R3.

Hexagonal View

The Hexagonal view (Figure 2) provides the users with an overview of one particular 2D projection of the dataset. In order to avoid overplotting when dealing with large datasets, the chemical compounds are aggregated into hexagonal bins. We opted for the hexagonal binning approach, first described by Carr et al. [7], as it evidences significant advantages for efficient data aggregation in comparison to other approaches. This is related to the low perimeter-to-area ratio of the hexagonal shape, which reduces the sampling bias. The ideal shape for that is a circle but it cannot be used for the full-coverage division of the plane. Hexagons are the most circular-shaped objects, enabling to construct such a division in an efficient way.

The Hexagonal view can be applied to any of the vector-based molecular representations described in Section 3.1. Each hexagon has an assigned opacity, which corresponds to the number of compounds aggregated within the hexagon. A higher opacity corresponds to a larger number of compounds inside the hexagon. In this way, the Hexagonal view depicts the distribution of the compounds yielded by the DR, so users can recognize areas with a high density of compounds. The color of a hexagon encodes the prevailing trend among its compounds for a selected feature, which is by default their bioactivity but can be switched to other molecular properties, including the trustworthiness of the projection (requirement R3).




Fig. 2: The Hexagonal view shows the density of the distribution of compounds in the 2D projection, which is encoded using into opacity, and the prevalence of the bioactivity of the compounds, which is encoded using color. The size of the hexagons can be adjusted.

The granularity of the Hexagonal view can be changed using a slider, which enables a more detailed view of the distribution of the compounds. In order to preserve the readability through the range of different granularity levels, the opacity is enhanced linearly with respect to the decrease of the hexagon size. Additionally, the Hexagonal view enables the user to select several hexagons at once. The compounds within the selected hexagons are then filtered through the other linked views (Detail view, 3D view, and Table view), where they can be explored in more detail.

Detail View

Upon selecting a subset of compounds in the Hexagonal view, the user can explore the selected data in the Detail view, depicted in Figure 3. In this view, the compounds are represented using a standard scatter plot, enhanced by a subtle overlay of the same hexagonal grid as in the Hexagonal view, which helps users keep the correspondence between the zoom level in these views. This is important because the Detail view displays only the selected hexagons zoomed in after the selection. To further enhance the link between the Hexagonal and Detail views, the corresponding hexagons are highlighted when the user hovers over them in any of these views. In order to perform a selection of individual compounds in this view, a lasso-shaped selector is supported. The selected compounds are then displayed in the 3D view and also highlighted in the Table view (Section 5.1.2).

As in the case of the Hexagonal view, the Detail view can be set to any of the vector-based molecular representations described in Section

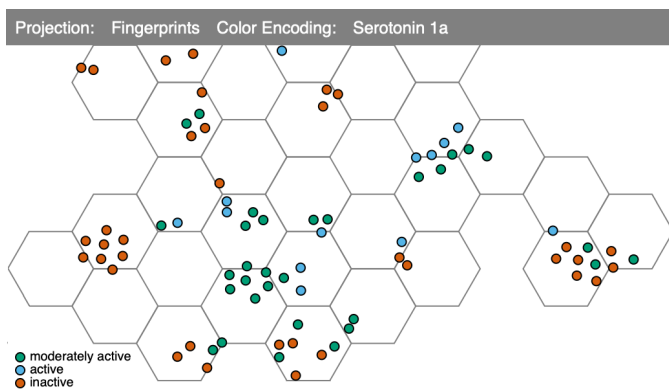


Fig. 3: The Detail view displays individual compounds as dots in a selected projection (ECFP fingerprints in this case). The color can be mapped into one of the features (activity towards Serotonin 1a receptor in this case).



Fig. 4: The Difference view showing the decomposition of the local neighborhoods from one projection to another one. Here, the reference projection *A* corresponds to ECFP fingerprints and its selected hexagons are highlighted in gray. Projection *B* corresponds to Daylight fingerprints (inner hexagons in black). The opacity of hexagons in *A* encodes the value of the metric used for the trustworthiness of projection (see Section 6).

3.1. In addition, all properties and features described in Section 3 can be color-encoded on points representing each compound. These features are selected from a drop-down menu, and their color encodings are chosen according to their type, i.e., quantitative, such as molecular weight, or categorical, such as bioactivity towards a target protein. Another quantitative property that can be used for color encoding corresponds to the correlation scores, which encode the trustworthiness of the DR projection of the compound (requirement R3). These scores were computed using Pearson and Kendall correlation, whose calculation is explained in the Supplementary Material.

Difference View

In order to support the task of comparing the outputs of different 2D projections, as stated in our requirement R2, we propose a novel view called *Difference view*, that combines and contrasts two selected 2D Hexagonal views, *A* and *B*. Initially it displays a hexagonal layout similar to that presented in the Hexagonal view, where the opacity of each hexagon encodes the computed correlation score of the trustworthiness of the projections *A* and *B* under study (requirement R3).

The Difference view adopts the hexagonal layout from one of the projections to be compared, which we denote as the reference projection. When we choose projection *A* as the reference and perform a selection operation in the Hexagonal view of *A*, we search for the positions of the compounds falling into this selection in projection *B*. These positions are encoded as inner smaller hexagons inscribed into the original hexagonal grid of the Difference view, as depicted in Figure 4. In other words, this view shows where the compounds from the selected hexagons in projection *A* are located in projection *B*. The size of inner hexagons corresponds to the number of compounds falling into the same hexagonal bin.

The primary purpose of this graph is to help illustrate which regions of the dataset preserve their neighborhoods from one vector-based molecular representation to another one when undergoing a DR procedure. In this way, the user can rapidly compare two different projections and assess the trustworthiness of the neighborhoods generated by the DR techniques. If the level of fragmentation from one projection to another is high, i.e., a segregation of the content of one hexagon into many small, fairly separated and scattered hexagons is observed, it can be inferred that these molecules behave differently according to the chosen molecular representation, thus they might be taken into consideration cautiously and inspected in more detail.




Fig. 5: The Table view. Each row corresponds to one compound and its selected features. The right side panel shows the statistical overview of the distribution of values across the dataset.

5.1.2 Table View

Besides from the vector-based molecular representations used in the DR projections and displayed in our 2D plot views, there are several other molecular features related to drug-likeness that should be taken into consideration when analyzing the compounds, as stated in requirement R4. ChemVA enables the users to explore such features, listed in Section 3.2, by means of a Table view which offers advanced interaction options. We adopted a well-established tool, published by Gratzl et al. [15], and its extension [13]. As these tools perfectly fit to our needs, we incorporate them to ChemVA. Further details about the broad range of interaction possibilities can be found in the original papers.

In addition to the list of compounds, the Table view provides the users with graphical elements in the form of juxtaposed bar charts and box plots in the right side panel. By default, the compounds in the table are logically grouped by their membership to hexagons in the Hexagonal view. These groups can be either expanded or compressed. When compressed, the table displays the box plots of the distribution of the values for each feature in the hexagon, as shown in Figure 5.

This view is interactively linked with the other visual components of ChemVA. When performing a selection in the 2D plot views, the corresponding compounds are automatically highlighted in the Table view. Conversely, when the user selects compounds in the Table view, the corresponding compounds are highlighted in the Detail view, displayed in the 3D view, and the corresponding hexagons are highlighted in the Hexagonal view.

5.1.3 3D View

The geometric similarity between selected compounds can be explored using our 3D view. The visual representation of atoms and bonds, as well as the coloring of the compounds, are based on standard representations used in molecular chemistry, with specific colors reserved for each chemical element. It also offers standard interactions, such as panning, zooming, and rotation of the displayed structures.




Fig. 6: 3D view with 66 aligned molecules. Common atoms and bonds are rendered more opaque than the less common components.




Fig. 7: Reducing the visual clutter visible in Figure 6 by (a) filtering out the non-common parts of compounds and by (b) changing the opacity of the whole structure.

Compound similarity can be better perceived when the compounds are structurally aligned in the view. To serve this purpose, we use a structural alignment functionality provided by the OpenBabel tool [51]. Further details about this functionality are provided in the Supplementary Material. Once molecules are aligned, the user should be able to easily identify their common parts, i.e., the subsets of atoms and bonds that are present in most of the selected compounds, as stated in requirement R5. In addition, we incorporate opacity modulation with respect to the frequency of occurrence of atoms and bonds. In other words, the opacity of atoms and bonds is calculated based on the number of atoms of the same type that are aligned to the same spatial position. As a consequence, common substructures are rendered more opaque. An example of such alignment can be seen in Figure 6.

Depending on the number of selected compounds and their similarities, the 3D visual representation can become fairly cluttered (Figure 6). To reduce this visual clutter, the 3D view offers hiding and showing hydrogen atoms and bonds on demand, changing the size of atoms and bonds, as it can be seen in Figure 7(a), and also changing the opacity of the whole structure. Figure 7(b) shows an example of using these functionalities on a cluttered subset of compounds.

In some cases, the user wants to focus only on the common parts of the structure. Therefore, we have also included functionality for filtering out atoms and bonds. Figure 7(a) shows how the visibility of the common substructure among 66 compounds is improved by using this feature. Finally, as the user might also want to analyze the non-common parts of the compounds, the 3D view provides an option to invert the opacity, so that the common part of the structure becomes more transparent than the rest.

5.2 Dimensionality Reduction

t-Distributed Stochastic Neighborhood Embedding (t-SNE) [77] is considered to be the state-of-the-art technique for dimensionality reduction, particularly for visualizing very high dimensional data. t-SNE maps neighboring data points to a lower-dimensional space, aiming to preserve the neighborhood relationships (*locality*). It applies the Student's t-distribution on the low-dimensional space aiming to tackle the *crowding problem*, which happens when many points clump together in the low-dimensional projection. These traits make t-SNE a suitable tool for visually exploring molecular spaces, where local neighborhoods are of utmost interest for analyzing compound similarity.

However, t-SNE has a major limitation, which is that it is a non-parametric technique. After testing *Parametric t-SNE* [76] and consistently failing to attain good quality projections, we developed our own parametric dimensionality reduction model. This model was constructed based on a feed-forward neural network, which was trained to learn the 2-dimensional coordinates of a previously computed t-SNE projection of the data. We trained four of these parametric models for each reference dataset used in ChemVA, one for each molecular representation. The performance of our parametric models was evaluated by measuring the Pearson correlation coefficient between the predicted coordinates and the actual coordinates, i.e., those obtained by means of the t-SNE projection. We provide details on the parameterization used for t-SNE and the parametric models in the Supplementary

Material. The t-SNE projections were computed using tools from Scikit-Learn [52], whereas the parametric models were built using Keras and Tensorflow [8].

5.3 Adding New Compounds

ChemVA provides an option to add new compounds to the dataset being studied in order to explore their features and compare them with those of other compounds. By means of this functionality, the expert can assess the potential of this newly added compound prior to extensive wet lab testing. The support for this important feature fulfills requirement R6.

A new compound is loaded to ChemVA by specifying its SMILES formula [83]. A back-end service of ChemVA calculates the molecular features listed in Section 3.2. After computing such molecular representations and features, ChemVA uses the parametric models (described in Section 5.2) to obtain the 2D coordinates of the compound in each projection. Finally, the newly added compound is displayed in all the supported views with yellow color to ease its identification. This color encoding prevails throughout the views, as it can be seen in Figure 8.

6 EVALUATION AND CASE STUDIES

The evaluation of ChemVA consisted of two main stages. The first stage was conducted by one domain expert involved in the functional design of the tool (see Section 4), and consisted of two case studies. The second stage was conducted by one visualization expert and two domain experts, who were not involved at any point during the design or development of ChemVA. This stage consisted of a session where the different views were qualitatively evaluated and general feedback about the usability and functionality of the tool was gathered afterwards.

6.1 First Stage - Case Studies

During the first stage of the evaluation, several additional functional requirements for the tool itself were identified by the domain expert, such as loading and storing selections, downloading the displayed 3D conformations, and highlighting the compounds in the Table view after clicking on them in the 3D view. These additional functional requirements were addressed in time for the case study evaluation.

ChemVA was tested on two case studies, which were built upon two different datasets retrieved from ChEMBL [14]. The first dataset was composed by merging ligands binding to the *Serotonin 1a receptor*¹ and *Dopamine D2 receptor*², whereas the second dataset comprised ligands to the *P-glycoprotein 1*³. We assigned a categorical label to each compound according to its experimentally measured *IC50* bioactivity value towards the target(s) under study. Compounds showing *IC50* values below 10 nM were labeled as *Active*; compounds between 10 and 1000 nM were labeled as *Moderately Active*, and those over 1000 nM were labeled as *Inactive*. The *serotonin-dopamine* dataset comprises 118 compounds, whereas the *P-glycoprotein* dataset contains 893 compounds.

6.1.1 Case Study 1: Analysis of Chemical Determinants of Activity Towards Serotonin and Dopamine Receptors

This case study is based on the *serotonin-dopamine* dataset, a set of compounds that have data about antagonistic activity against the serotonin 5HT1A and dopamine D2 receptor. These receptors of neurotransmitters are targets of many psychoactive pharmaceuticals, i.e., antidepressants, antipsychotics, and anxiolytics. In the field of psychopharmacology, small molecules with different activity towards several receptors are used to alleviate side-effects [48, 67].

The goal of the study was to find chemical determinants of biological activity towards serotonin and dopamine receptors. According to the domain expert, similar compounds could be easily found in the Detail view in all four projections, based on their proximity. After identifying groups of potentially similar compounds, the domain expert searched for groups of compounds that were *active* towards both receptors and

that had desirable pharmacological properties, such as following the Lipinski's RO5 and having a high QED score. This exploratory search was done on the Hexagonal view, by adjusting the size of the hexagons to match the observed groups of compounds in the different projections.

Afterwards, the domain expert used the Kendall and Pearson correlation color encoding in order to observe whether the projections were trustworthy, and thus the hexagons being considered would effectively group similar compounds. A set of hexagons was preselected and the Table view was used, in which the domain expert analyzed the distributions of all drug-likeness features of the selected compounds using the summary information displayed in the column headers. As a result, the expert identified a hexagon that grouped active compounds towards both receptors with desirable drug-likeness features.

The domain expert selected those compounds and performed an alignment in the 3D view, finding that their structures were very similar. These compounds were also contrasted to inactive compounds within the same hexagon using the 3D view. This allowed the domain expert to gain insight into the relevant substructures of a potential new drug candidate. At this point, the domain expert highlighted how easy it was to find and visualize the common 3D structure of a group of compounds.

Finally, the structures of active compounds towards both receptors were downloaded from ChemVA. Five new compounds were created and loaded to ChemVA, as shown in Figure 8. They were compared to the downloaded structures using molecular docking. This process is described in the Supplementary Material. The domain expert found out that the new structures effectively bound to a known antagonist binding pocket and three out of the five newly created structures showed slightly better binding energies than structures from the dataset used for their design, which were already highly active. This shows that ChemVA could be effectively used for drug design purposes, leading to newly designed compounds that have the desired qualities and bioactivity profiles.

6.1.2 Case Study 2: Analysis of Structural Determinants of P-glycoprotein Inhibitors

This case study was based on the *P-glycoprotein* dataset, consisting of small molecules with inhibitory activity against human P-glycoprotein. This protein is exclusively overexpressed in cells of many cancer types, causing multidrug resistance of these cells and thus impacting on the performance of chemotherapy treatments. It also affects the effectiveness of many drugs by altering their ADME-Tox properties. [68]. P-glycoprotein has proven to bind to many structurally dissimilar substrates, thus an interesting use case for our tool was to find structural determinants of a good P-glycoprotein inhibitor and to compare known active ligands towards this target.

According to the domain expert, the P-glycoprotein dataset contains a vast number of diverse compounds that interact with the target protein. The goal of this case study was to find chemical determinants of compounds with a very high logP value, which are scarce in the dataset. In order to achieve this goal, the domain expert used the Table view to filter compounds with logP values higher than 6.75 and then sorted them based on their logP value and their correlation scores, which measure the trustworthiness of the molecular representation currently being projected onto the 2D plot views.

The domain expert used the filtering options provided by the Table view in order to find a set of compounds complying with all the requested criteria, i.e., high logP value, activity towards the P-glycoprotein and following Lipinski's RO5. Afterwards, chemically similar compounds were selected from this subset and further analyzed using the Hexagonal view. Hexagons containing active compounds were selected and thoroughly explored using the Detail view, in which logP was used for color encoding. Finally, the domain expert used the 3D view to find common substructures among the selected compounds and successfully identified chemical determinants of bioactivity. This was achieved by analyzing structural patterns related to lipophilicity, indicated by high logP values. This workflow is depicted in Figure 9.

¹https://www.ebi.ac.uk/chembl/target_report_card/CHEMBL214/

²https://www.ebi.ac.uk/chembl/target_report_card/CHEMBL217/

³https://www.ebi.ac.uk/chembl/target_report_card/CHEMBL4302/




Fig. 8: The newly designed compounds were projected by ChemVA near the selected dataset compounds in the Daylight projection (top), occupying the same hexagons, and slightly further from them in the ECFP Fingerprints projection (bottom). The Difference view on the bottom right illustrates the redistribution of the selected regions. a) Hexagonal view, b-d-f) Detail view, c) 3D view, e) Difference view, g) Table view.

6.2 Second Stage - Qualitative Evaluation by External Experts

In the second stage we introduced ChemVA to two domain experts from the Loschmidt Laboratories at Masaryk University who evaluated our tool in terms of functionality and user-experience. We also introduced our newly proposed Difference graph to one visualization expert (see Acknowledgments), who has vast experience in DR techniques. None of the experts were involved at any stage during the design and development of ChemVA.

First, we conducted a brief introduction to the layout and the proposed visualization methods. This stage took approximately thirty minutes for each of the evaluations with the domain experts, after which they were able to use the tool without any further assistance. While both of the domain experts highlighted the intuitiveness of the Hexagon and 3D views, they also pointed out the need for an appropriate documentation describing other views, such as our novel Difference view, in order to support usability. According to the domain experts, ChemVA allows to easily compare numerous properties, which is an uncommon but very useful feature in tools for virtual screening. Both of the domain experts highlighted the usefulness of having different projections, and the ability to focus on a specific source of data using the correlation scores and comparing them by means of the Difference view. According to the visualization expert, the Difference view does a good job by enabling the user to visualize how the two mappings, hence the two feature spaces, are correlated in terms of grouping the same set of molecules together. Both domain experts enumerated several examples from their own ongoing research activities which they claimed could be enhanced by ChemVA, in terms of time and effort. For instance, a drug design task aimed at identifying *on- and off-target* drug responses, or the analysis of 4300 FDA-approved drugs for usability in the treatment of COVID-19, which is currently carried out using cumbersome scripting and manual feature selection.

7 RESULTS AND DISCUSSION

In this section we summarize some final remarks and feedback provided by the domain experts who conducted both stages of the evaluation of ChemVA. We discuss the results of both case studies, as well as strengths and limitations of our tool. As shown in Section 6.1, the goals for both case studies were effectively fulfilled using ChemVA, which suggests that the tool has the potential to be widely adopted and used by medicinal chemists. This claim is also supported by the results of the qualitative evaluation performed by external experts, described in Section 6.2. Our tool allowed an intuitive comparison of sets of chemically similar compounds. In contrast to other tools, which only allow to sort and filter compounds by a set of properties, ChemVA enabled a comprehensive analysis of the data by taking into consideration multiple molecular representations and levels of granularity. In the case studies, this was aided by specific features of our tool, such as grouping compounds in hexagons and allowing to adjust their size.

In the first case study, the Detail view assisted the domain expert in distinguishing similar compounds based on proximity in all four projections (R1 & R2). Furthermore, the colors and shapes presented in the Hexagonal and Detail views allowed the domain expert to find a specific hexagon very quickly (R1).

The task of analyzing compound similarity could be thoroughly tested by means of different projections provided by ChemVA (R2). Furthermore, during the two case studies it was possible to examine these projections in terms of their trustworthiness, using the provided correlation scores (R3), and to compare them by means of the Difference view (R2). These two features provided by ChemVA aided the domain expert in the first case study in order to assess whether the projections were trustworthy and thus hexagons were effectively grouping similar compounds.

The evaluation also showed the usefulness of the 3D view to discover and analyze similar 3D structures among the selected compounds

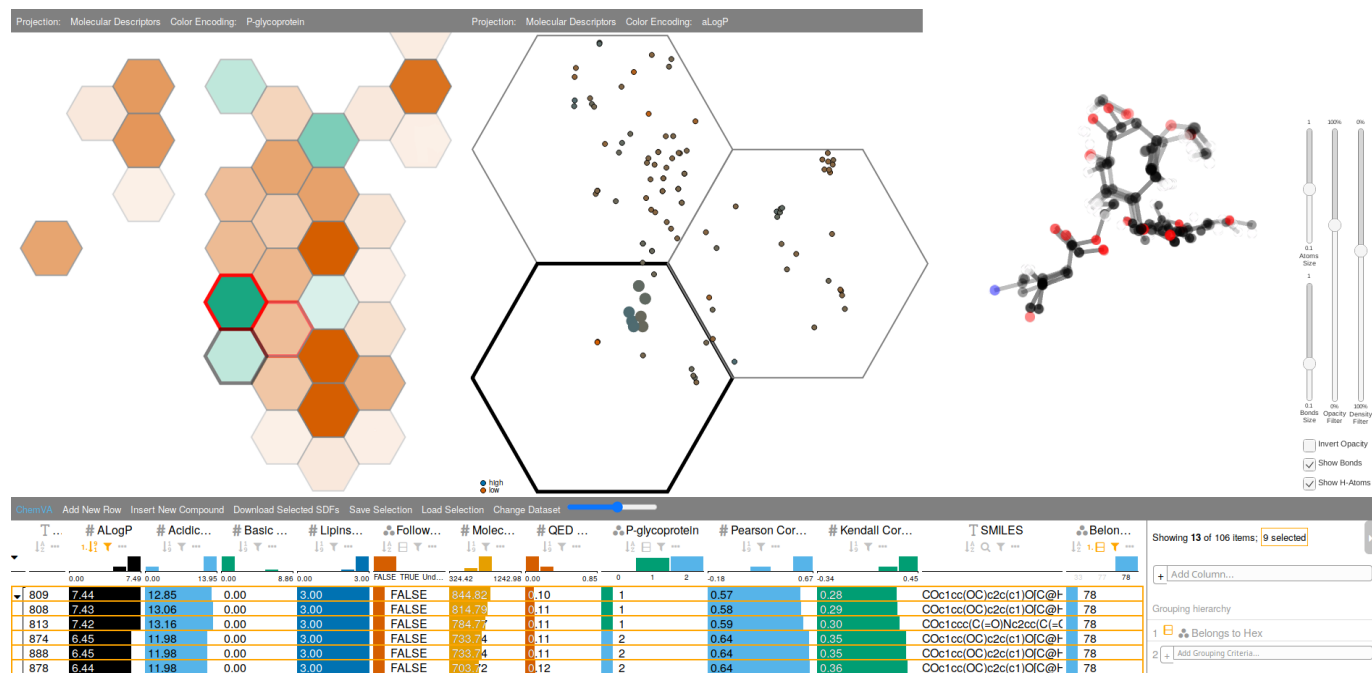


Fig. 9: Illustration of the exploratory process in the P-glycoprotein dataset, displaying only a subset of interesting molecules in the Detail view and selecting those with high logP values.

(R5). Examining the chemical determinants for bioactivity in both case studies was straightforward by means of the molecular alignment function provided by ChemVA, according to the domain expert. Every subset of molecules could be quickly aligned on demand, in contrast to other chemical software tools, where this task constitutes a strenuous process and it is often restricted to molecules above a certain threshold of structural similarity [16]. Moreover, structural patterns could be further explored using the *opacity filter* and *invert opacity option*, which allow highlighting the different parts of the compounds.

The coordination of the plots with the Table view was also found useful by the domain experts (R4), offering fast access to summary information in its column headers. The Table view proved to be effective for exploring a large dataset, such as *P-glycoprotein*, enabling the search for rare feature values through filtering and hierarchical grouping during the second case study. Other functionalities, such as showing SMILES formulas for copying, the option for exporting 3D structures, or searching compounds by SMILES formula, made ChemVA a flexible tool, easily interoperable with other cheminformatics tools.

One limitation identified during the first case study is that new compounds were not always projected near to structurally similar compounds. This functionality performed consistently better on some molecular representations than the others, as it can be seen in Figure 8. This might be related to the fact that the dataset serotonin-dopamine used in the experiment is small (118 compounds), thus the neural-based parametric model trained for projecting new compounds is not able to generalize well for every molecular representation. Another limitation we identified is that not all compounds have available 3D structures, which may limit the usage of the 3D view with some datasets.

ChemVA proved to be useful both as an exploration and visualization tool, and it also helped the domain expert with the design and evaluation of new compounds (R6) during the first case study. As discussed in Section 6.2, ChemVA was also praised as a useful and innovative tool by the domain experts during the external evaluation.

We also identified potential directions for future extensions of the tool, accounting both for user experience and functionality. The external domain experts suggested expanding the current functionality of the tool by enabling the user to load and customize their own data. One of the domain experts also suggested to enable adding new compounds by drawing them in a sketcher. Other possible directions include providing options for substructure search or to enhance navigation by providing

breadcrumbs and saving analytical snapshots. Lastly, it is worth noticing that although our tool is domain-specific, ChemVA has been developed by following techniques and strategies that could be applied to other types of data, including non-chemical data. For instance, our newly proposed Difference view could assist in the process of visually analyzing the preservation of clusters and neighborhoods for most clustering methods. Also, the analysis of trustworthiness of DR projections could be applicable to any type of data undergoing a DR procedure. The idea behind the 3D view, which allows for overlapping common molecular substructures, could be applied to other graph-based data or even text.

8 CONCLUSIONS

In this paper, we presented ChemVA, a novel tool for visual analysis of chemical compounds, which is especially focused on evaluating molecular similarity for virtual screening. Our tool proposes a set of coordinated views that support visual exploration and comparison of 2D projections, which are obtained from applying dimensionality reduction on different—and ideally, complementary—molecular representations. Our case studies, conducted by a medicinal chemist, confirmed that ChemVA addresses the functional requirements and provided appropriate support for his analysis on two different datasets. The qualitative evaluation conducted by three other external experts revealed the potential of our tool for its adoption in the domain, as well as the usefulness of our newly proposed Difference view. The evaluation process allowed us to identify potential extensions to ChemVA, which will steer our future work on the tool.

ACKNOWLEDGMENTS

The presented work has been supported by DFG-GACR research project no. GC18-18647J, by CONICET research grant PIP 112-2017-0100829, by UNS research grants PGI 24/N042 and PGI 24/N048, and by ANPCyT (Argentina) research grant PICT-2017-1246. We thank Sérgio M. Marques and Ondřej Vávra from the Loschmidt Laboratories, and Michaël Aupetit from the Qatar Computing Research Institute QCRI for evaluating ChemVA and for their insightful comments and suggestions on the features of the tool.

SUPPLEMENTARY MATERIAL

In this Supplementary Material, we provide the readers with detailed information about data preprocessing performed for ChemVA and computational times, as well as the results of the molecular docking performed for the Case study. Although this document contains important information for reproducing our approach and the experiments conducted during the evaluation of the tool, its contents are not crucial for understanding the visualization concepts used in ChemVA, which constitute the core of our proposal.

DATA PREPROCESSING

In this section we describe the computation of the vector-based molecular representations and the equations used for computing the correlation scores. We also provide details on the parameterization of our DR method, parametric models and molecular alignment back-end service. The information presented in this section is relevant for reproducibility purposes.

Vector-based Molecular Representations

In order to provide domain experts with diverse information about the ligands in our datasets, we computed four different molecular representations, which we then used as sources of data for our 2-dimensional projections. For each dataset, we computed radius 2 ECFPs and Daylight fingerprints, both 1024 bits long, using the *Chem* and *AllChem* packages of RDKit [36]. We also computed 0D, 1D, and 2D molecular descriptors using Mordred [46], obtaining a total of 1613 descriptors. We removed all descriptors that had more than 10% *NaN* values and replaced all remaining *NaN* with a fixed value (maximum value per descriptor), which yielded 1454 molecular descriptors. Finally, we used a pretrained Mol2Vec model [23], which computed 300-dimensional embeddings based on the SMILES formulas of the compounds in the datasets. The computation of ECFPs, Daylight fingerprints and molecular descriptors took less than an hour, and it was performed in a 32-core cluster with 12GB memory. This preprocessing stage was performed offline, which means that it was done once and then the results were stored for future use.

Features for Drug-Likeness

The features related to drug likeness are displayed in the Table view and can be color encoded onto the Detail and Hexagonal views as well. The values for such features were obtained using a publicly available REST API by ChEMBL [14]. The collected values did not undergo any further preprocessing steps.

t-SNE Projections and Parametric Models

We computed a t-SNE projection for each of the vector-based molecular representations, for both of the datasets used in the case study. The t-SNE projections were computed using the *manifold* class from Scikit-Learn [52], and the parameters used were selected by grid search. The tested perplexity values varied from 5 to 10% of the total amount of compounds in the dataset. The same parameterization on each dataset was used to fit t-SNE for each of the vector-based molecular representations. Fitting each t-SNE projection took around 20 minutes on a 32-core cluster with 12GB memory. A summary of the parameterization used for t-SNE is provided in Table 1.

Table 1: Summary of the parameters used to fit the t-SNE projections of each dataset.

Dataset	P-glycoprotein	Serotonin-Dopamine
Perplexity	45	5~10
Maximum No. epochs	10000	10000
No. epochs without progress	500~1000	1000

After fitting each t-SNE projection, we trained one parametric model to learn each of such projections. The parametric models were built using Keras and Tensorflow [8]. The parameterization of these models

is summarized in Tables 2 and 3. The computational time invested on training these models was of approximately 30 minutes per model, on a 32-core cluster with 12GB memory. Both the fitting of the t-SNE projections and the training of the parametric models consisted of offline tasks.

Correlation Scores for Measuring the Trustworthiness of Projections

t-SNE applies two different distributions to map data points to a lower-dimensional space in an effort to preserve neighborhoods according to a specific optimization criterion. As a result, pairwise distances among the compounds in the low-dimensional space could be distorted with respect to distances in the high-dimensional space, and those differences could mislead the domain expert during the virtual screening process. For this reason, as a means to assess the trustworthiness of the projections, we computed two different *correlation scores* for each of the compounds.

The first score, namely r , was computed by calculating the cosine distance of each compound k to the rest of the compounds in the high-dimensional space and then measuring the Pearson correlation (see Equation 1) among these distances and the ones in the low-dimensional space. Pearson correlation has been previously used by Strickert et al in MSR [69], which has also been applied to chemical compounds. The second score, namely τ , was computed by generating two ranks for each compound k , one in the origin space and another one in the mapped space, which sort all the compounds according to the cosine distance to k . Afterwards, the score was computed by comparing both ranks using the Kendall rank correlation (see Equation 2).

Both correlation scores were computed for each compound individually, and once for each of the four vector-based molecular representations used in ChemVA, which yields four different views on which the trustworthiness of each projection can be assessed. In addition, the *difference graph* uses the Kendall rank correlation scores as the default encoding. These scores were computed using tools in the Python Data Science suite SciPy [28] and Scikit-Learn [52].

$$r = \frac{n(\sum xy) - (\sum x)(\sum y)}{\sqrt{[n\sum x^2 - (\sum x)^2][n\sum y^2 - (\sum y)^2]}} \quad (1)$$

$$\tau = \frac{n_{\text{concordants}} - n_{\text{discordants}}}{n(n-1)/2} \quad (2)$$

Molecular Alignment and 3D Conformations

ChemVA supports a 3D view, which allows the user to conduct a detailed analysis of the molecular structures under study, and to visually explore common geometrical patterns among compounds. The 3D conformations of compounds were retrieved from PubChem [30] and computed using OpenBabel [51]. The computed 3D conformations were obtained using *Ghemical* force field energy minimization, force field cleanup (500 cycles) and slow rotor search. The computation of these 3D conformations consisted of an offline task. It took approximately between 2 and 20 minutes for each compound, depending on the complexity of its molecular structure, and it was performed on a 32-core cluster with 12GB memory.

The molecular alignment function involves two steps. The first step consists of finding the *maximum common substructure* (MCS) among the selected compounds, which is achieved using RDKit [36]. The second step consists of aligning all the selected molecules with respect to the found MCS, and it is done using the *obfit* function provided by OpenBabel [51]. The molecular alignment task is performed *on demand*: it is an online task, which means that it is performed upon the selection of a set of compounds in the Detail view in our back-end servers. The response time of this web service is of approximately one second, tested on random selections from 5 to 30 compounds.

Adding New Compounds

The addition of new compounds to the Detail view is an online task, performed *on demand* in our back-end servers. The parametric models

Table 2: Summary of the parameters used to build and train the four parametric models for *P-glycoprotein* dataset.

Parameters	ECFPs	Daylight fps	Molecular descriptors	Molecular embeddings
# Nodes per hidden layer	100 / 10 / 2	100 / 10 / 2	200 / 20 / 2	100 / 10 / 2
Activation fn	relu	relu	relu	relu
Dropout coefficients	0.25 / 0.15 / 0.1	0.25 / 0.15 / 0.1	0.25 / 0.15 / 0.1	0.25 / 0.15 / 0.1
Early stopping patience	70	130	120	50
Early stopping delta	0.005	0.005	0.005	0.005
Learning rate (Adam opt)	0.0001	0.0001	0.0001	0.0001

Table 3: Summary of the parameters used to build and train the four parametric models for *serotonin-dopamine* dataset.

Parameters	ECFPs	Daylight fps	Molecular descriptors	Molecular embeddings
# Nodes per hidden layer	50 / 10 / 2	50 / 10 / 2	200 / 50 / 2	50 / 10 / 2
Activation fn	relu	relu	tanh	sigmoid
Dropout coefficients	0.25 / 0.15 / 0.1	0.25 / 0.15 / 0.1	0.25 / 0.15 / 0.1	0.25 / 0.15 / 0.1
Early stopping patience	70	70	100	70
Early stopping delta	0.005	0.005	0.005	0.005
Learning rate (Adam opt)	0.0001	0.0001	0.0001	0.0001

are previously loaded, which reduces significantly the response time: approximately four seconds per compound.

CASE STUDY 1 - MOLECULAR DOCKING

This section provides details on the two datasets used in the case studies, as well as on the results of the molecular docking study performed by the domain expert in the context of our case study 1 (see Section 6 of the paper).

Dataset Selection and Preprocessing

The two datasets used in the case studies, namely *serotonin-dopamine*⁴⁵ and *P-glycoprotein*⁶, were built upon two different datasets retrieved from ChEMBL [14]. We performed a preliminary selection step of the SMILES formulas on each of the datasets, keeping only those compounds that could be loaded using RDKit [36]. Details on the amount of compounds belonging to each class are provided in Tables 4 and 5. The compounds were labeled according to their experimentally measured *IC50* bioactivity value towards the targets under study.

Table 4: Details on the composition per class of the *P-glycoprotein* dataset.

# compounds per class	Class	P-glycoprotein
	Active	42
Moderately active	178	
Inactive	673	
Total	893	

Table 5: Details on the composition per class of the *serotonin-dopamine* dataset.

# compounds per class	Class	Serotonin 1a	Dopamine D2
	Active	14	5
Moderately active	42	28	
Inactive	62	85	
Total	118	118	

Preparation of Receptor Structures

The structure of human 5HT1A receptor was constructed using homology modeling using the *I-TASSER* web server using default set-

⁴https://www.ebi.ac.uk/chembl/target_report_card/CHEMBL214/

⁵https://www.ebi.ac.uk/chembl/target_report_card/CHEMBL217/

⁶https://www.ebi.ac.uk/chembl/target_report_card/CHEMBL4302/

tings [88]. The model was constructed using homologous 5HT receptors and monoamine receptors with resolved structures. Model 1 was selected for the docking study with a TM-score of 0.52 ± 0.15 , indicating a well-predicted fold [92]. The model was aligned to the homologous structure of 5HT1B (PDB ID 4IAR [79]) and showed a nearly perfect overlap. The structure of the dopamine D2 receptor was downloaded from PDB [4] (PDB ID 6CM4 [81]). Hydrogens, ions, and waters were removed using PyMOL [64]. The protonation state of titratable groups, as well as the orientation of Asn, Gln, and His groups, were corrected by the H++ web server using pH 7.5 [1].

Preparation of Ligand Structures

Based on the three structures selected in ChemVA, five new ligand structures were designed using Avogadro [18]. Molecular docking aimed to compare binding energies and binding modes of the three dataset structures and those of the five designed structures. The molecular docking was done using Autodock Vina [75]. Autodock atom types and Gasteiger charges were added to the two receptors and the ligands using MGLTools [47, 63]. The docking grid was selected to be a $30 \times 30 \times 30 \text{ \AA}$ cube. For the 5HT1A receptor, the cube was centered on the bound ligand in the aligned structure of 5HT1B. For the dopamine D2 receptor, the cube was centered on the bound ligand in the receptor structure.

Results and Discussion

All of the seven structures successfully bound to the binding pockets of the two receptors, suggesting a plausible binding mode was found. Three of the five designed structures showed a lower value binding energy, indicating a stronger binding. The binding energies of the best binding modes are shown in Table 6, and the SMILES formulas for each of the designed compounds are provided in Table 7.

In this comparison, designed compounds **2**, **4**, and **5** exhibited higher predicted affinities to both receptors. These predicted results do not necessarily mean higher inhibitory activity, given that such bioactivity can be influenced by other factors such as flexibility of the receptor proteins.

In addition, we provide supplementary files comprising the results from the I-TASSER web server and the input files and results from the Autodock Vina procedure.

REFERENCES

- [1] R. Anandakrishnan, B. Aguilar, and A. V. Onufriev. H++ 3.0: automating p k prediction and the preparation of biomolecular structures for

Table 6: Predicted binding affinities towards the *serotonin 5HT1A* receptor (*5HT1A*) and *dopamine D2 receptor* (*dopD2*) of both selected dataset compounds and designed compounds, measured in *kcal/mol*. Rows in bold show the compounds exhibiting the highest predicted affinities to both receptors.

Compound	Affinity to 5HT1A	Affinity to dopD2
Compound 45	-8.0	-8.7
Compound 79	-8.0	-8.3
Compound 117	-9.9	-11.0
Designed compound 1	-8.4	-11.5
Designed compound 2	-10.5	-11.9
Designed compound 3	-9.9	-10.8
Designed compound 4	-10.3	-11.2
Designed compound 5	-10.2	-11.2

Table 7: SMILES formulas of the designed compounds. Rows in bold show the compounds exhibiting the highest predicted affinities to both receptors.

Designed Compound	SMILES formula
1	<chem>Fc1ccc(cc1)c2ccc(CNCC3CCc4cccc4C3)c2</chem>
2	<chem>Fc1ccc(cc1)c2ccc(CNCC3CCc4ccc(C)c(C)cc4C3)c2</chem>
3	<chem>N(Cc1ccc(c1)c2cc(F)ccc2F)CC3CC4=C(C=CC=C4)O3</chem>
4	<chem>N(Cc1ccc(c1)c2ccc(F)cc2)C(N(C)C)C3CCC4=C(C=CC=C4)O3</chem>
5	<chem>Fc1ccc(cc1)c2ccc(c2)C(N)CC3CCc4cccc4O3</chem>

atomistic molecular modeling and simulations. *Nucleic Acids Research*, 40(W1):W537–W541, 2012.

[2] M. Aupetit. Visualizing distortions and recovering topology in continuous projection techniques. *Neurocomputing*, 70(7-9):1304–1330, 2007. doi: 10.1016/j.neucom.2006.11.018

[3] F. Barbosa and D. Horvath. Molecular similarity and property similarity. *Current Topics in Medicinal Chemistry*, 4(6):589–600, 2004. doi: 10.2174/1568026043451186

[4] H. Berman, J. Westbrook, Z. Feng, G. Gilliland, T. Bhat, H. Weissig, I. Shindyalov, P. Bourne, P. Rose, A. Prlic, et al. Rcsb protein data bank: Structural biology views for basic and applied research. *Nucleic Acids Res*, 28:235–242, 2000.

[5] D. Borland, W. Wang, J. Zhang, J. Shrestha, and D. Gotz. Selection bias tracking and detailed subset comparison for high-dimensional data. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):429–439, 2019. doi: 10.1109/TVCG.2019.2934209

[6] M. Bostock, V. Ogievetsky, and J. Heer. D³ data-driven documents. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2301–2309, 2011. doi: 10.1109/TVCG.2011.185

[7] D. B. Carr, R. J. Littlefield, W. Nicholson, and J. Littlefield. Scatterplot matrix techniques for large n. *Journal of the American Statistical Association*, 82(398):424–436, 1987. doi: 10.2307/2289444

[8] F. Chollet. Keras, 2015. <https://github.com/fchollet/keras>, online July 2020.

[9] M. A. Cox and T. F. Cox. Multidimensional scaling. In *Handbook of Data Visualization*, pp. 315–347. Springer, 2008. doi: 10.1007/978-3-540-33037-0_14

[10] R. Cutura, S. Holzer, M. Aupetit, and M. Sedlmair. VisCoDeR: A tool for visually comparing dimensionality reduction algorithms. In *Esann*, 2018.

[11] Daylight Chemical Information Systems, Inc. Daylight fingerprints, 2019. <http://www.daylight.com/>, online April 2020.

[12] R. A. Fisher. The use of multiple measurements in taxonomic problems. *Annals of Eugenics*, 7(2):179–188, 1936. doi: 10.1111/j.1469-1809.1936.tb02137.x

[13] K. Furmanová, S. Gratzl, H. Stitz, T. Zichner, M. Jarešová, A. Lex, and M. Streit. Taggle: Combining overview and details in tabular data visualizations. *Information Visualization*, 19(2):114–136, 2020. doi: 10.1177/1473871619878085

[14] A. Gaulton, A. Hersey, M. Nowotka, A. P. Bento, J. Chambers, D. Mendez, P. Mutowo, F. Atkinson, L. J. Bellis, E. Cibrián-Uhalte, et al. The ChEMBL database in 2017. *Nucleic Acids Research*, 45(D1):D945–d954, 2017. doi: 10.1093/nar/gkw1074

[15] S. Gratzl, A. Lex, N. Gehlenborg, H. Pfister, and M. Streit. LineUp: Visual

analysis of multi-attribute rankings. *IEEE Transactions on Visualization and Computer Graphics (InfoVis '13)*, 19(12):2277–2286, 2013. doi: 10.1109/tvcg.2013.173

[16] M. Gütlein, A. Karwath, and S. Kramer. CheS-Mapper 2.0 for visual validation of (Q)SAR models. *Journal of Cheminformatics*, 6(1), 2014. doi: 10.1186/s13321-014-0041-7

[17] R. M. Hanson, J. Prilusky, Z. Renjian, T. Nakane, and J. L. Sussman. JSmol and the next-generation web-based representation of 3d molecular structure as applied to proteopedia. *Israel Journal of Chemistry*, 53(3-4):207–216, 2013. doi: 10.1002/ijch.201300024

[18] M. D. Hanwell, D. E. Curtis, D. C. Lonie, T. Vandermeersch, E. Zurek, and G. R. Hutchison. Avogadro: an advanced semantic chemical editor, visualization, and analysis platform. *Journal of cheminformatics*, 4(1):17, 2012.

[19] R. P. Hertzberg and A. J. Pope. High-throughput screening: new technology for the 21st century. *Current Opinion in Chemical Biology*, 4(4):445–451, 2000. doi: 10.1016/S1367-5931(00)00110-1

[20] T. Hoffmann and M. Gastreich. The next level in chemical space navigation: going far beyond enumerable compound libraries. *Drug Discovery Today*, 24(5):1148–1156, 2019. doi: 10.1016/j.drudis.2019.02.013

[21] W. Humphrey, A. Dalke, and K. Schulten. VMD: Visual molecular dynamics. *Journal of Molecular Graphics*, 14(1):33–38, 1996. doi: 10.1016/0263-7855(96)00018-5

[22] Y. G. Ichihara, M. Okabe, K. Iga, Y. Tanaka, K. Musha, and K. Ito. Color universal design: the selection of four easily distinguishable colors for all color vision types. In *Color Imaging XIII: Processing, Hardcopy, and Applications*, vol. 6807, p. 680700. International Society for Optics and Photonics, 2008. doi: 10.1117/12.765420

[23] S. Jaeger, S. Fulle, and S. Turk. Mol2vec: Unsupervised machine learning approach with chemical intuition. *Journal of Chemical Information and Modeling*, 58(1):27–35, 2018. doi: 10.1021/acs.jcim.7b00616

[24] A. P. Janssen, S. H. Grimm, R. H. Wijdeven, E. B. Lenselink, J. Neeffjes, C. A. van Boeckel, G. J. van Westen, and M. van der Stelt. Drug discovery maps, a machine learning model that visualizes and predicts kinase-inhibitor interaction landscapes. *Journal of Chemical Information and Modeling*, 59(3):1221–1229, 2018. doi: 10.1021/acs.jcim.8b00640

[25] Jmol: an open-source java viewer for chemical structures in 3d, 2009. www.jmol.org, online January 2020.

[26] M. A. Johnson and G. M. Maggiora. *Concepts and applications of molecular similarity*. Wiley, 1990. doi: 10.1002/jcc.540130415

[27] I. M. Johnstone and D. M. Titterton. Statistical challenges of high-dimensional data. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 367(1906):4237–4253, 2009. doi: 10.1098/rsta.2009.0159

[28] E. Jones, T. Oliphant, P. Peterson, et al. SciPy: Open source scientific tools for Python, 2001–. [Online; accessed].

[29] Khronos. WebGL: OpenGL ES for the Web. <https://www.khronos.org/webgl/>, online April 2020, 2011.

[30] S. Kim, P. A. Thiessen, E. E. Bolton, J. Chen, G. Fu, A. Gindulyte, L. Han, J. He, S. He, B. A. Shoemaker, et al. Pubchem substance and compound databases. *Nucleic Acids Research*, 44(D1):D1202–d1213, 2016. doi: 10.1093/nar/gkv951

[31] T. Kohonen. The self-organizing map. *Proceedings of the IEEE*, 78(9):1464–1480, 1990. doi: 10.1109/5.58325

[32] B. Kovalerchuk. *Visual knowledge discovery and machine learning*, vol. 144. Springer, 2018. doi: 10.1007/978-3-319-73040-0

[33] B. Kozlíková, M. Krone, M. Falk, N. Lindow, M. Baaden, D. Baum, I. Viola, J. Parulek, and H.-C. Hege. Visualization of biomolecular structures: State of the art revisited. *Computer Graphics Forum*, 36(8):178–204, 2016. doi: 10.1111/cgf.13072

[34] M. A. Kramer. Nonlinear principal component analysis using autoassociative neural networks. *AIChE Journal*, 37(2):233–243, 1991. doi: 10.1002/aic.690370209

[35] E. Krieger and G. Vriend. YASARA View—molecular graphics for all devices—from smartphones to workstations. *Bioinformatics*, 30(20):2981–2982, 2014. doi: 10.1093/bioinformatics/btu426

[36] G. Landrum. Rdkit: open-source cheminformatics <http://www.rdkit.org>, 2016.

[37] R. A. Laskowski and M. B. Swindells. LigPlot+: Multiple ligand-protein interaction diagrams for drug discovery. *Journal of Chemical Information and Modeling*, 51(10):2778–2786, 2011. doi: 10.1021/ci200227u

[38] R. Lewis, R. Guha, T. Korcsmaros, and A. Bender. Synergy maps: exploring compound combinations using network-based visualization. *Journal*

- of *Cheminformatics*, 7(1), 2015. doi: 10.1186/s13321-015-0090-6
- [39] E. Lionta, G. Spyrou, D. K. Vassiliadis, and Z. Cournia. Structure-based virtual screening for drug discovery: Principles, applications and recent advances. *Current Topics in Medicinal Chemistry*, 14(16):1923–1938, 2014. doi: 10.2174/1568026614666140929124445
- [40] S. Liu, D. Maljovec, B. Wang, P.-T. Bremer, and V. Pascucci. Visualizing high-dimensional data: Advances in the past decade. *IEEE Transactions on Visualization and Computer Graphics*, 23(3):1249–1268, 2016. doi: 10.2312/eurovisstar.20151115
- [41] R. Macarron, M. N. Banks, D. Bojanic, D. J. Burns, D. A. Cirovic, T. Garyantes, D. V. Green, R. P. Hertzberg, W. P. Janzen, J. W. Paslay, et al. Impact of high-throughput screening in biomedical research. *Nature Reviews Drug Discovery*, 10(3):188–195, 2011. doi: 10.1038/nrd3368
- [42] G. M. Maggiora and V. Shanmugasundaram. Molecular similarity measures. In *Cheminformatics*, pp. 1–50. Springer, 2004. doi: 10.1385/1-59259-802-1:001
- [43] L. McInnes, J. Healy, and J. Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.
- [44] M. R. Min, H. Guo, and D. Shen. Parametric t-distributed stochastic exemplar-centered embedding. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 477–493. Springer, 2018. doi: 10.1007/978-3-030-10925-7_29
- [45] K. R. Moon, D. van Dijk, Z. Wang, D. Burkhardt, W. S. Chen, A. van den Elzen, M. J. Hirn, R. R. Coifman, N. B. Ivanova, G. Wolf, et al. Visualizing transitions and structure for high dimensional data exploration. *bioRxiv*, 2017. doi: 10.1101/120378
- [46] H. Moriwaki, Y.-S. Tian, N. Kawashita, and T. Takagi. Mordred: a molecular descriptor calculator. *Journal of Cheminformatics*, 10(1):4, 2018.
- [47] G. M. Morris, R. Huey, W. Lindstrom, M. F. Sanner, R. K. Belew, D. S. Goodsell, and A. J. Olson. Autodock4 and autodocktools4: Automated docking with selective receptor flexibility. *Journal of Computational Chemistry*, 30(16):2785–2791, 2009.
- [48] H. Nasrallah. Atypical antipsychotic-induced metabolic side effects: insights from receptor-binding profiles. *Molecular Psychiatry*, 13(1):27–35, 2008. doi: 10.1038/sj.mp.4002066
- [49] J. J. Naveja and J. L. Medina-Franco. Finding constellations in chemical space through core analysis. *Frontiers in Chemistry*, 7, 2019. doi: 10.3389/fchem.2019.00510
- [50] L. G. Nonato and M. Aupetit. Multidimensional projection for visual analytics: Linking techniques with distortions, tasks, and layout enrichment. *IEEE Transactions on Visualization and Computer Graphics*, 25(8):2650–2673, 2019. doi: 10.1109/TVCG.2018.2846735
- [51] N. M. O’Boyle, M. Banck, C. A. James, C. Morley, T. Vandermeersch, and G. R. Hutchison. Open babel: An open chemical toolbox. *Journal of Cheminformatics*, 3(1), 2011. doi: 10.1186/1758-2946-3-33
- [52] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [53] J. Peltonen and Z. Lin. Information retrieval approach to meta-visualization. *Machine Learning*, 99(2):189–229, 2015. doi: 10.1007/s10994-014-5464-x
- [54] E. F. Pettersen, T. D. Goddard, C. C. Huang, G. S. Couch, D. M. Greenblatt, E. C. Meng, and T. E. Ferrin. UCSF Chimera—A visualization system for exploratory research and analysis. *Journal of Computational Chemistry*, 25(13):1605–1612, 2004. doi: 10.1002/jcc.20084
- [55] D. Probst and J.-L. Reymond. Exploring drugbank in virtual reality chemical space. *Journal of Chemical Information and Modeling*, 58(9):1731–1735, 2018. doi: 10.1021/acs.jcim.8b00402
- [56] D. Probst and J.-L. Reymond. Visualization of very large high-dimensional data sets as minimum spanning trees. *Journal of Cheminformatics*, 12(1):1–13, 2020. doi: 10.1186/s13321-020-0416-x
- [57] P. E. Rauber, A. X. Falcão, and A. C. Telea. Visualizing time-dependent data using dynamic t-SNE. In E. Bertini, N. Elmqvist, and T. Wischgoll, eds., *EuroVis 2016 - Short Papers*. The Eurographics Association, 2016. doi: 10.2312/eurovisshort.20161164
- [58] N. Rego and D. Koes. 3Dmol.js: molecular visualization with WebGL. *Bioinformatics*, 31(8):1322–1324, 2015. doi: 10.1093/bioinformatics/btu829
- [59] D. Rogers and M. Hahn. Extended-connectivity fingerprints. *Journal of Chemical Information and Modeling*, 50(5):742–754, 2010. doi: 10.1021/ci100050t
- [60] A. Ronacher. Flask: a lightweight WSGI web application framework, 2009. <https://palletsprojects.com/p/flask/>, online April 2020.
- [61] D. Sacha, L. Zhang, M. Sedlmair, J. A. Lee, J. Peltonen, D. Weiskopf, S. C. North, and D. A. Keim. Visual interaction with dimensionality reduction: A structured literature analysis. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):241–250, 2016. doi: 10.1109/tvcg.2016.2598495
- [62] T. Sander, J. Freyss, M. von Korff, and C. Rufener. DataWarrior: An open-source program for chemistry aware data visualization and analysis. *Journal of Chemical Information and Modeling*, 55(2):460–473, 2015. doi: 10.1021/ci500588j
- [63] M. F. Sanner, A. J. Olson, and J.-C. Spehner. Reduced surface: an efficient way to compute molecular surfaces. *Biopolymers*, 38(3):305–320, 1996.
- [64] L. Schrödinger. The PyMOL molecular graphics system, version 1.8, 2015.
- [65] M. Sedlmair, M. Brehmer, S. Ingram, and T. Munzner. Dimensionality reduction in the wild: Gaps and guidance. *The University of British Columbia, Tech. Rep.*, 2012.
- [66] M. Sedlmair, T. Munzner, and M. Tory. Empirical guidance on scatterplot and dimension reduction technique choices. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2634–2643, 2013. doi: 10.1109/tvcg.2013.153
- [67] S. Sifakis, D. Tzachanis, M. Samara, and G. Papazisis. Antipsychotic drugs: from receptor-binding profiles to metabolic side effects. *Current Neuropharmacology*, 16(8):1210–1223, 2018. doi: 10.2174/1570159x15666170630163616
- [68] K. M. R. Srivalli and P. Lakshmi. Overview of p-glycoprotein inhibitors: a rational outlook. *Brazilian Journal of Pharmaceutical Sciences*, 48(3):353–367, 2012. doi: 10.1590/s1984-82502012000300002
- [69] M. Strickert, A. J. Soto, and G. E. Vazquez. Adaptive matrix distances aiming at optimum regression subspaces. In *ESANN 2010, European Symposium on Artificial Neural Networks*, pp. 93–98, 2010.
- [70] J. Tang, J. Liu, M. Zhang, and Q. Mei. Visualizing large-scale and high-dimensional data. In *Proceedings of the 25th International Conference on World Wide Web*, pp. 287–297, 2016. doi: 10.1145/2872427.2883041
- [71] Y. Tanrikulu, B. Krüger, and E. Proschak. The holistic integration of virtual screening in drug discovery. *Drug Discovery Today*, 18(7-8):358–364, 2013. doi: 10.1016/j.drudis.2013.01.007
- [72] U. Technologies. Unity Engine, 2005. <https://unity.com>, online March 2020.
- [73] S. Tilkov and S. Vinoski. Node.js: Using javascript to build high-performance network programs. *IEEE Internet Computing*, 14(6):80–83, 2010. doi: 10.1109/MIC.2010.145
- [74] R. Todeschini and V. Consonni. *Handbook of molecular descriptors*, vol. 11. John Wiley & Sons, 2010. doi: 10.1002/9783527613106
- [75] O. Trott and A. J. Olson. Autodock vina: improving the speed and accuracy of docking with a new scoring function, efficient optimization, and multithreading. *Journal of Computational Chemistry*, 31(2):455–461, 2010.
- [76] L. Van Der Maaten. Learning a parametric embedding by preserving local structure. In *Artificial Intelligence and Statistics*, pp. 384–391, 2009.
- [77] L. Van Der Maaten and G. Hinton. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008.
- [78] L. Van Der Maaten, E. Postma, and J. Van den Herik. Dimensionality reduction: a comparative review. *Journal of Machine Learning Research*, 10(66-71):13, 2009.
- [79] C. Wang, Y. Jiang, J. Ma, H. Wu, D. Wacker, V. Katritch, G. W. Han, W. Liu, X.-P. Huang, E. Vardy, et al. Structural basis for molecular recognition at serotonin receptors. *Science*, 340(6132):610–614, 2013.
- [80] J. Wang, X. Liu, and H.-W. Shen. High-dimensional data analysis with subspace comparison using matrix visualization. *Information Visualization*, 18(1):94–109, 2019. doi: 10.1177/1473871617733996
- [81] S. Wang, T. Che, A. Levit, B. K. Shoichet, D. Wacker, and B. L. Roth. Structure of the d2 dopamine receptor bound to the atypical antipsychotic drug risperidone. *Nature*, 555(7695):269–273, 2018.
- [82] Y. C. Wang, Q. Zhang, F. Lin, C. K. Goh, and H. S. Seah. Polarviz: a discriminating visualization and visual analytics tool for high-dimensional data. *The Visual Computer*, 35(11):1567–1582, 2019. doi: 10.1007/s00371-018-1558-y
- [83] D. Weininger. Smiles. 3. depict. graphical depiction of chemical structures. *J. Chem. Inf. Comput. Sci.*, 30(3):237–243, 1990. doi: 10.1021/

ci00067a005

- [84] S. Wold, K. Esbensen, and P. Geladi. Principal component analysis. *Chemometrics and Intelligent Laboratory Systems*, 2(1-3):37–52, 1987. doi: 10.1016/0169-7439(87)80084-9
- [85] J. Wu, J. Wang, H. Xiao, and J. Ling. Visualization of high dimensional turbulence simulation data using t-SNE. In *19th AIAA Non-Deterministic Approaches Conference*, 2017. doi: 10.2514/6.2017-1770
- [86] X. Xu, T. Liang, J. Zhu, D. Zheng, and T. Sun. Review of classical dimensionality reduction and sample selection methods for large-scale data processing. *Neurocomputing*, 328:5–15, 2019. doi: 10.1016/j.neucom.2018.02.100
- [87] J. Xue, H. Zhang, and K. Dana. Deep texture manifold for ground terrain recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 558–567, 2018. doi: 10.1109/CVPR.2018.00065
- [88] J. Yang, R. Yan, A. Roy, D. Xu, J. Poisson, and Y. Zhang. The I-TASSER Suite: protein structure and function prediction. *Nature Methods*, 12(1):7, 2015.
- [89] A. Yoshimori, T. Tanoue, and J. Bajorath. Integrating the structure–activity relationship matrix method with molecular grid maps and activity landscape models for medicinal chemistry applications. *ACS Omega*, 4(4):7061–7069, 2019. doi: 10.1021/acsomega.9b00595
- [90] W. Yu and A. D. MacKerell. *Computer-Aided Drug Design Methods*, pp. 85–106. Springer New York, New York, NY, 2017. doi: 10.1007/978-1-4939-6634-9_5
- [91] G. G. Zanabria, L. G. Nonato, and E. Gomez-Nieto. iStar (i*): An interactive star coordinates approach for high-dimensional data exploration. *Computers & Graphics*, 60:107–118, 2016. doi: 10.1016/j.cag.2016.08.007
- [92] Y. Zhang and J. Skolnick. Scoring function for automated assessment of protein structure template quality. *Proteins: Structure, Function, and Bioinformatics*, 57(4):702–710, 2004.
- [93] W. Zhu, Z. Webb, X. Han, K. Mao, W. Sun, and J. Romagnoli. Generic process visualization using parametric t-SNE. *IFAC-PapersOnLine*, 51(18):803–808, 2018. doi: 10.1016/j.ifacol.2018.09.262